

Inference and Learning for Rigid-Body Models of Manipulation

by

Nima Fazeli

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

~~August 2019~~ [September 2019]

© Massachusetts Institute of Technology 2019. All rights reserved.

Signature redacted

Author

.....
Department of Mechanical Engineering
August 1, 2019

Signature redacted

Certified by




.....
Alberto Rodriguez
Associate Professor
Thesis Supervisor

Signature redacted

Accepted by

Nicolas Hadjiconstantinou
Department Graduate Officer



Inference and Learning for Rigid-Body Models of Manipulation

by

Nima Fazeli

Submitted to the Department of Mechanical Engineering
on August 1, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

In this thesis, we explore a spectrum of inference and modeling approaches for robotic manipulation. Particularly, we investigate the broad class of rigid-bodies undergoing frictional interactions.

We begin by deriving a contact-implicit system identification formulation for articulated rigid-bodies. Assuming we have a physical model of the system, the objective is to derive system parameters and contact forces for articulated rigid-bodies without enumerating and inferring contact formations. We then ground this approach by investigating the fidelity of rigid-body contact models and their identification. We evaluate the fidelity of the contact models by empirically studying their predictive performance and parameter identification properties in a planar impact task.

Next, we address one approach to augmenting these contact models with data. The objective here is to improve model fidelity through an optimization of model parameters and residual error learning for systems with prior physics models.

We conclude the thesis by building models from data for tasks with rich latent structure and no prior physics models. Here, the objective is to learn data-efficient hierarchical models of physics that incorporate force and tactile sensory modalities and are amenable to inference, controls, and planning.

Thesis Supervisor: Alberto Rodriguez

Title: Associate Professor

Acknowledgments

I am eternally grateful to my advisor, Prof. Alberto Rodriguez. His guidance and advice has been a shining light on my path through research. I have had the great fortune of having an outstanding advisor and life-long friend. Alberto's care for his students is par none and he has taught me the true meaning of being an academic advisor.

I am very grateful for my lab mates, for all the intellectual conversations, for all the good times we spent together in and out of the lab, and for all the bonds we'll share for a life-time. I am particularly grateful to Nikhil Chavan-Dafle, François R. Hogan, and Rachel M. Holladay for putting up with my nonsense, I truly cherish our friendship.

I sincerely thank my committee, Professors Neville Hogan and Leslie P. Kaelbling for their time and invaluable feedback. Professor Hogan was one of the main reasons why I had wanted to pursue my PhD at MIT, having seen one of his talks at DSCC in the early 2010s. He was not taking on students the year I joined. Professor Kaelbling taught me how to think about algorithms and learning systems, her insights and methods will stay with me for the rest of my career.

Last but not least, I am eternally grateful to my family for all their support, love, and understanding. Being apart from my parents and sister for this many years has been difficult, yet I've always felt their warmth from half a world away. Finally, none of this would have been possible without my partner, Saba Jahanbin. She has been my rock and grounding. Her love has never wavered, through thick and thin, thank you for putting up with me for all these years.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction	21
1.1	Challenges	22
1.2	Importance	23
1.3	Contributions	23
1.3.1	Contact-implicit identification of articulated rigid-bodies	23
1.3.2	Empirical contact model predictive performance evaluation	24
1.3.3	Sample-efficient residual learner models for point contacts	25
1.3.4	Hierarchical model learning for complex manipulation skills	25
1.4	Thesis Structure	26
2	Contact-Implicit System Identification for Articulated Rigid-Bodies	27
2.1	Goal	27
2.2	Introduction	28
2.3	Problem Formulation	30
2.3.1	Time-stepping complementarity formulation for simulation of rigid-bodies undergoing frictional contact	30
2.3.2	Contact-implicit system identification	34
2.3.3	Identifiability analysis	39
2.4	Experiments - Simulation	43
2.5	Experiments - Real World	50
2.6	Related Work	53
2.7	Discussion	56

3	Empirical Contact Model Predictive Performance Evaluation	59
3.1	Goal	59
3.2	Introduction	60
3.3	Problem Formulation	60
3.3.1	Fundamental constraints of contact	62
3.3.2	Rigid contact models studied in this paper	65
3.3.3	Model identification and evaluation	67
3.3.4	Experimental data	68
3.4	Results	71
3.5	Related Work	78
3.6	Discussion	81
4	Sample-efficient Residual Learner Contact Models	83
4.1	Goal	83
4.2	Introduction	84
4.3	Problem Formulation	85
4.3.1	Review of the Energy Ellipse	85
4.3.2	Learning rigid-body contact models	87
4.3.3	Uncertainty propagation	90
4.3.4	Illustrative Example:	92
4.4	Results	94
4.5	Related Work	101
4.6	Discussion	102
5	Topic Models of Manipulation	103
5.1	Goal	103
5.2	Introduction	104
5.3	Formulation	105
5.3.1	Model realization	108
5.3.2	Model learning	110
5.3.3	Topic model analogy	112

5.4	Experiments	114
5.4.1	Real-world setup	114
5.4.2	Simulation setup	116
5.4.3	Task specifications and data collection	117
5.4.4	Evaluation metric	118
5.5	Results	119
5.5.1	Performance on evaluation metric	119
5.5.2	Simulation	119
5.5.3	Real-world	120
5.5.4	Model learning	122
5.6	Related Work	126
5.7	Discussion	128
6	Closing Thoughts	133
6.1	Why model-based or compositional approaches?	133
6.2	On simulation to real-world transfer (sim2real)	135
6.3	State-estimation and perception are hard	136
6.4	The action-perception loop	137

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

2-1	Schematic of the frictional interaction between a two link manipulator and a disk and the disk with the ground.	28
2-2	Is the trajectory of the rod in the figure indicative of the dynamic system that governs its motion? Note the two key events: When the left end of the rod makes contact and sticks to the ground, and the subsequent contact when the right end of the rod contacts the ground and slides. [80] used the example of a falling rod to introduce a time-stepping complementarity scheme for contact resolution that has become one of the standard techniques for simulating frictional contact. In this paper we look at the same formulation and similar examples from the perspective of identification.	29
2-3	Distance to contact. For the example of a planar object falling on a horizontal surface, we consider the distances ϕ_n and ϕ_t between the closest point to the ground P and the origin O of a fixed reference frame. The point P is parametrized by the configuration of the object $\mathbf{q} = (x, y, \theta)$ and the angle β	31
2-4	A disk transitions from free flight to rolling contact. The solid line denotes the true trajectory of the disk, and the dots denote observations corrupted by sensor noise. Given an observation window of length m we want to estimate Θ and $\{\mathbf{c}^1 \dots \mathbf{c}^m\}$ for the duration of the window.	36
2-5	Uncertainty in the pose measurement can lead to configurations that do not lie on the contact manifold. Projection algorithms or relaxation in the constraints can be used to account for this uncertainty	39

2-6	Block in free-fall.	44
2-7	Ellipse in free fall.	45
2-8	Rimless wheel rolling/sliding downhill	46
2-9	a) LCP simulation trace of a falling block. b) LCP simulation trace of a falling ellipse. c) LCP simulation trace of a falling rolling rimless wheel for the cases of high friction (top) and low friction (bottom) with the ground.	47
2-10	Dependence of the motion of the rimless wheel with its mass.	48
2-11	Normal contact force of the rimless wheel spokes as a function of time.	49
2-12	Overlay of estimated and measured contact force between the manipulator and disk. The contact forces and inertial parameters are estimated concurrently. The contact forces are estimated to within approximately 20%. The discrepancy is partially attributed to the short time horizon of the interaction providing limited persistence of excitation.	51
2-13	Experimental dropping setup. A six axis industrial robot latches magnetically on a part, gives it some initial velocity and orientation and drops it. A motion capture system and a F/T sensor capture the falling motion at high frequency.	51
2-14	Example of a falling trajectory of a planar ellipse in the experimental dropping arena. The figure shows a very sparse sub-sample of the trajectory recorded by the motion tracking system.	52
2-15	Reconstruction of the contact force for a impact of an ellipse.	53
3-1	Planar object making contact with a horizontal surface. The contact point is denoted by C	62

3-2	The Energy Ellipse: The ellipse is defined in the contact frame where the ordinate and abscissa are the normal and tangential impulses, respectively. The interior of the ellipse demarcates the set of impulses that conserve kinetic energy before and after contact. Contact impulses above the <i>Line of Maximum Compression</i> prevent penetration between the bodies making contact. Impulses on the <i>Line of Sticking</i> result in zero tangential velocity during contact (sticking contact). The impulses in the shaded region satisfy the constraints of rigid-body contact models.	64
3-3	Regions of the admissible contact space covered by each model. Note the similarity between the models and the limited coverage of the admissible space.	66
3-4	Experimental setup: Autonomous data collection for planar impact data. A planar object is constrained to move between two lubricated vertical glass planes. The robot uses a directional magnetic latch to pick up the object across the glass and provide control over position and velocity of the object. The Vicon motion capture system is used to track the states of the object.	69
3-5	The configurations of the object as a function of time during an experimental run. θ is in radians but has been scaled by the radius of gyration of the object.	70
3-6	Example test of object energy during the first two bounces. The net energy of the object is approximately constant between two impact events. The net energy drops after each impact event.	70
3-7	The left panel shows the topography of the cost function plotted as a function of the parameters for the Whittaker model where the summation is over 150 drops and lighter colors denote better objective values. The right panel shows the convergence behaviour of the model parameters and objective function values for 3 randomly selected subsets of data, plotted as a function of number of data in the subset k	72

3-8	Estimated probability density functions (PDFs) of the ℓ_2 -norm error in post-contact velocity of the center of mass.	74
3-9	The performance of the model depicted as a function of the normal and tangential components of the pre-contact momentum of the system measured at the contact point. Deep blue regions represent high accuracy, whereas yellow regions predict high error regions.	75
3-10	Identified values of (μ, ϵ) for each drop, where a) shows the values color-coded by the ability of the identified pair to predict the impulse, and b) showing the distribution of the parameters of the data-set.	76
3-11	a) Example of $\mu = 0$, this impact is characterized by the point of contact being located directly under the center of mass of the object. In this situation the Energy ellipse is axis aligned to the normal impulse direction. b) Example of a measured impulse outside the convex set of predictions made by the model due to "spin-reversal".	78
4-1	a) planar object making contact with a fixed horizontal surface, b) Energy Ellipse.	87
4-2	The new Energy Ellipsoid due to the addition of the contact wrench to account for deformations and finite contact time.	89
4-3	Empirically measured trajectory of the first two contacts of a dice roll, overlaid with samples from the posterior probability distribution from the stochastic simulation. The samples are generated by the stochastic process used to propagate uncertainty.	93
4-4	Posterior distribution over horizontal distance of the center of mass of the dice normalized by the dice side length, from the stochastic simulator. Plots show initial configuration and the two post contact configurations.	94
4-5	The estimated probability density function of the ℓ_2 norm error in predicted post contact velocity for 450 data samples used to train the models. a) Class I, b) Class II.	95

4-7	An example dice roll rendered from empirically collected data. To generate the plot, we sample the trajectory measured from motion tracking uniformly and plot the trace of the outline of the block for each point.	96
4-8	Error distributions over 750 drops normalized by the length of the side of the dice.	97
4-9	On the left is an example of a difficult to predict roll, where 2 temporally close consecutive impacts occur and a small error significantly impacts future contact events, where as on the right a single clean impact occurs that can be predicted more reliably	100
4-10	Distribution over the possible horizontal resting positions of the dice for an easy to predict and difficult to predict roll normalized by the length of the side of the dice. The distribution is over-layed with the expected output of the analytical model.	100
5-1	Model Learning in the Wild: A Robot learning the mechanics of interacting with a Jenga tower. Building useful temporal abstractions over these mechanics has powerful implications for prediction, planning, and controls. For example, inferring high-level block behaviors such as “free” or “stuck” early can simplify planning considerably. These behaviors are governed by micro-frictional interactions and unobservable geometric difference in blocks. Fortunately, tactile/force sensing provides rich contextual clues that help build abstractions that encode the effect of these mechanics.	104
5-2	Example: Here, a robot pusher is attempting to move the purple block. The block in (a) the block is very heavy and in (b) it is light enough to be displaced. How can the robot build useful abstractions such as “not movable” with physics predictions?	106

5-3 Plate Notation for the Hierarchical Physics Model: We can interpret the model as a filter of length D over N number of trajectories. K denotes the number of abstractions (unknown a priori). Details of model construction are provided in the text. 107

5-4 Robot Setup: Physical setup consisting of the robot, Jenga tower, Intel RealSense D415 camera, and ATI Gamma Force/Torque sensor (mounted at the wrist). (B) The machine intelligence architecture with the learned physics model. 115

5-5 Jenga setup in simulation and the baseline comparisons: A) The simulation setup is designed to emulate the real-world implementation. B) The learning curve of the different approaches with confidence intervals evaluated over 10 attempts. C) A visual depiction of the structure of the mixture of regressions (MOR) and the proposed approach (HMA). 116

5-6 Summary statistics for Exploration and Learned Physics. Table comparing the performance of the robot using the exploration strategy and the learned model. 121

5-7 Learned Intuitive Physics: (A) Overlay of the analytical friction cone and predicted forces given the current measurements. The friction coefficient between the finger material (PLA) and wood is between 0.35 and 0.5, here we use 0.42 as an approximation. (B) The normal force applied to the tower as a function of the height of the tower. . . 122

5-8 Inference using the learned representation: Evolution of the beliefs of the robot as it interacts with the tower. (A) For a block that is stuck. (B) For the block that moves easily. 124

5-9 Controlled block pushing: The robot selects the point on the block and the appropriate angle to push with such that it realigns the block with the goal configuration. Here, the block is beginning to rotate counter clockwise and is starting to move out of the tower. The robot selects a point close to the edge of the block and pushes it back in towards to tower center. We convert angles to normalized distances by scaling with the radius of gyration of the block, 0.023 meters. We have exaggerated the block translation to illustrate the fine grain details of motion. 125

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

2.1	2D Numerical Simulation and Identification Results	48
3.1	Average-best parameter identified values and the percentage a model was chosen as the best and worst model	73
3.2	Fraction of drops that lie within the predictive range of each model .	78

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

Manipulation: An agents control over its interaction with the environment through selective contact [61].

This definition of manipulation implies a notion of predictive ability and understanding in our agents (here, robots). One effective way of generating this understanding is through models that allow robots to make inferences about the world.

More specifically, manipulation in uncertain environments can benefit from an explicit understanding of contact. Robots rely on *a priori* models of objects and their physical interactions as a means of estimating the state of their environment, planning towards their goals, and controlling their actions.

When available, physics-based analytical *a priori* models attempt to make predictions of the state of the environment given robot actions. These models are, inevitably, deficient or defective somehow: in some cases it is not cost-effective to build accurate models; in others the lack of observability of the governing factors in the interaction renders it impossible. Further, parameteric models of interactions need to be optimized for the instance of the world in situ – i.e. system identification that explicitly accounts for contact.

When unavailable, these models need to be built from experience. These physics models must be expressive enough to account for the complex and hybrid nature of contact. Finding representations that are sufficiently descriptive of the world without prohibitive sample-complexity remains an open challenge in robotics.

Whether analytical, data-driven, or on the spectrum between, the explicit understanding of the algebra between motions and forces is central to monitoring, control, and planning of manipulation skills. In this thesis, our objective is to develop means to building this understanding in the form of models for physical interactions.

1.1 Challenges

Contact is central to manipulation. As robots touch their environment, they change it. The characteristic challenge in manipulation is that robots have to reason about and cause this change in partially unknown environments using noisy and incomplete sensory information. Some aspects of this challenge are:

- **Modeling:** Contact is a complex phenomenon, it is governed by difficult to observe micro-interactions and is sensitive to initial conditions. Researchers employ a variety of modeling simplifications to build computationally efficient models useful for robotics. Many of these simplifications induce some amount of error in predictions. Further, manipulation tasks exist for which we do not have physical models. Consider the game of Jenga, where the block interactions are governed by unobservable micro-interactions between the blocks.
- **Multi-modality:** Contact induces hybrid dynamics (e.g. in contact or not). Further, most interesting manipulation tasks are under-actuated, i.e. fewer control dimensions than degrees of freedom. Reasoning over hybrid physical interactions is still an open problem in robotics. For example, how do we infer the system parameters of a rigid-body making and breaking contact without breaking up the trajectory of the object into non-contact and contact interactions.
- **Partial observability:** Partial observability has many facets in manipulation. One simple example is when reaching into a shelf, a robot may occlude its visual perception and has to rely on tactile feedback. A richer example is when a robot turns the handle of a locked door. The interaction is fully observable from a

traditional state perspective, but the fact that the door is locked is captured by a (potentially unknown) latent state. Partial observability introduced complex challenges in reasoning for contact.

There are clear connections between these aspects. For instance, we'd like models that are able to reproduce the multi-modal nature of contact while providing a mechanism to make inferences over partially observable latent states. We'd also like planning and control algorithms that are able to reason over the hybrid nature of contact to produce carefully executed manipulation actions.

1.2 Importance

Manipulation is key to realizing the promise of robotics: a solution to some of society's biggest challenges – including patient care, industrial automation, and disaster response. Central to these challenges is a robot's ability to control its environment through selective contact. Developing models and inference algorithms that facilitate manipulation is a concrete step towards this promise.

1.3 Contributions

This thesis provides four contributions to the state-of-the-art in inference and modeling for robotic manipulation:

1.3.1 Contact-implicit identification of articulated rigid-bodies

Identifying object physical properties such as inertia, friction, and restitution can provide important information for manipulation planning and control.

In this section, we derive an algorithm to infer the system properties and contact forces for articulated rigid-bodies making and breaking frictional contact. We assume that the frictional interaction is closely approximated as rigid and use Linear Complementarity contact models ([79, 4]).

As input, the algorithm takes a time history of system trajectories (positions and velocities) together with the geometry of the bodies and i) derives the set of identifiable system parameters ii) solves for the identifiable set of system parameters (contact and inertia terms) together with the contact forces. This approach is contact-implicit, i.e. it accounts for the sequence of contacts events in the optimization.

With this approach, we partially address the first two challenges in sec. 1.1 by providing a mechanism to handle the hybrid dynamics of contact implicitly for a specific class of objects.

1.3.2 Empirical contact model predictive performance evaluation

In the previous section, we assumed that the models we use for contact (in robotics) are sufficiently accurate to predict outcomes of physical interactions. How much should we trust these predictions?

In this section, we study the empirical performance of rigid-body contact models in a planar impact task. More specifically, we study the discrepancy between analytical models and measured data on a planar impact data-set and provide i) an empirical comparison and evaluation of 6 common rigid-body contact models used in robotics and computer graphics on a planar impact data-set; ii) a systematic approach to model parameter identification and a discussion on variance of parameter identification; and iii) an empirical upper-bound on performance of these contacts models for the planar impact data-set and a discussion on some of the notable issues and short-comings of these models.

Here, we investigate the influence of the first challenge in sec. 1.1 by providing the first empirical comparison and evaluation of several common rigid-body contact models used in robotics.

1.3.3 Sample-efficient residual learner models for point contacts

If a contact model is not sufficiently accurate for a task, how can we improve its predictive performance using experience?

In this section, we leverage the mathematical construction of the contact models in the previous section to formulate data-efficient residual error models. The objective here is to correct the erroneous predictions of the analytical contact models with the residual learner in the contact impulse space where we exploit a number of invariances in the input state-space. We further explore uncertainty propagation through these models using a Gaussian particle filtering technique.

In this section, we partially address the second and third challenges in sec. 1.1 by building on analytical models with data and reasoning over multi-modal uncertainty.

1.3.4 Hierarchical model learning for complex manipulation skills

What if the robot is faced with a physical interaction for which it does not have a pre-defined model? How can it acquire models with mechanisms to capture abstractions over physical interactions that are useful for inference, planning, and controls?

In this section, we study the problem of building hierarchical models of physics for complex manipulation skills. We leverage the semantically rich information provided by tactile feedback in combination with visual feedback to build useful abstractions over physical interactions with meaningful latent structure. We implement the approach on a robot learning the mechanics of its interaction with a Jenga tower. In this problem, the latent structure is in the types of behaviors a block may have (such as stuck or not), where tactile feedback provides important context to these behaviors.

In this section, we partially address all three challenges in sec. 1.1 by building on models capable of expressing hybrid and multi-modal contact mechanics induced by unobservable micro-interactions such as those of blocks in a Jenga tower.

1.4 Thesis Structure

The two common threads connecting the thesis topics are modeling and inference. While these threads serve to maintain coherency across chapters, each chapter is self-contained. As such, each chapter address background and related work within it's own scope and context. A chapter is dedicated to each of the aforementioned contributions in the sequence discussed.

This thesis draws heavily from the authors prior published work. In chapter 1, we review and extend [30, 29] by introducing additional methodologies for treating uncertainty in states. In chapter 2, we review and extend [28, 31] by providing the a refined derivation of the identification procedure. In chapter 3, we review and extend [33] by providing a mechanism to propagate uncertainty. In chapter 4, we review and extend [27] by providing a grounding and analogy to topic models used in natural language processing and a formal and general derivation of the model used in the original publication (which can be seen as a special case).

We conclude the thesis with a chapter on closing thoughts. This chapter is dedicated to some of the lessons learned and open problems in manipulation.

Chapter 2

Contact-Implicit System

Identification for Articulated Rigid-Bodies

2.1 Goal

Consider the system depicted in Fig. 2-1. Our objective is to estimate:

- System inertial properties,
- Contact forces and properties (friction and restitution).

We have access to:

- Joint torques,
- States (positions and/or velocities),
- Geometry (object shapes).

We assume rigid-body motion and rigid-body contact. Further, we assume that the contact parameters are fixed for the duration of observations – in the next chapter we investigate a state-dependent parameterization of the contact models. More formally, we present a constructive approach to deriving the identifiable set of parameters and a means of estimating this set given a sequence of observations.




Figure 2-1: Schematic of the frictional interaction between a two link manipulator and a disk and the disk with the ground.

2.2 Introduction

Central to successful autonomous manipulation is the ability to predict outcomes of actions. Contact is a central phenomenon in manipulation that must be handled carefully. To this end, robots can leverage *a priori* models of objects and the environment when available.

These models are low-dimensional representations of physics, parameterized by variables that can be adapted to the environment. Key to models useful for manipulation is an explicit understanding of the algebra between motions, forces, and inertia at contact. These models not only reason about free motion governed by inertial properties, but also contact properties such as friction and restitution.

In this chapter, we present a constructive formulation for the derivation of the identifiable set of parameters/variables for these models together with an estimation of their values. We're inspired by our own effective ability to make sense of our environment through contact. It only takes us a small push to a cup of coffee to estimate how full it is, and a quick glance to a bouncing ball to gauge its stiffness. Similarly, this work aims for robots to harness known laws of physical interaction to make sense of observed motions and/or forces, and as a result gain a better understanding of their environment and themselves.

The key challenge, and focus of this chapter, is in deriving a formulation suitable




Figure 2-2: Is the trajectory of the rod in the figure indicative of the dynamic system that governs its motion? Note the two key events: When the left end of the rod makes contact and sticks to the ground, and the subsequent contact when the right end of the rod contacts the ground and slides. [80] used the example of a falling rod to introduce a time-stepping complementarity scheme for contact resolution that has become one of the standard techniques for simulating frictional contact. In this paper we look at the same formulation and similar examples from the perspective of identification.

for system identification that can handle the complexity of contact dynamics with unknown and intermittent reaction forces due to frictional contact. Such a formulation provides the basis of a systematic approach for identification in a broader set of contact interactions. For example, consider the rod trajectory depicted in Fig. 2-2. The formulation needs to handle the free flight motion and sticking/sliding frictional interactions in a unified framework (as opposed to trying to break the trajectory into contact modes).

To this end, we build on the time-stepping linear complementarity approach to simulation of rigid-body frictional interactions [79, 4]. Through a re-parameterization of the dynamics, we derive the identifiable set of (inertia and contact) parameters for a given system. Further, we derive an optimization program leveraging this re-parameterization that computes an estimation of the identifiable set for a window of observed interactions.

2.3 Problem Formulation

In this section, we derive the identification formulation. To this end, we first review the time-stepping approach to simulating rigid-bodies undergoing frictional interaction. We then build on this approach to formulate the system identification optimization program. We end this section by discussing parameter identifiability.

2.3.1 Time-stepping complementarity formulation for simulation of rigid-bodies undergoing frictional contact

One approach to simulating rigid-body dynamical systems undergoing frictional interactions is to use the time-stepping complementarity formulation proposed by Stewart and Trinkle [79]. In this section, we provide a brief review of the approach within the context of the system identification formulation. For further details, see [79, 4, 66, 29].

Consider the systems depicted in Fig. 2-1. During the contact-free phase of the motion, we may write the equations governing the dynamics as:

$$\mathbf{M}(\mathbf{q})\frac{d\mathbf{v}}{dt} + \mathbf{C}(\mathbf{q}, \mathbf{v})\mathbf{v} = \mathbf{g}(\mathbf{q}) + \mathbf{f}_{\text{ext}}(\mathbf{q}) \quad (2.1)$$

These equations are also known as the Manipulator equations. Here, \mathbf{q} and \mathbf{v} respectively represent the joint configurations and velocities. The left side encodes the motion of the system where:

- $\mathbf{M}(\mathbf{q})$ denotes the (positive-definite) joint space inertia matrix.
- $\mathbf{C}(\mathbf{q}, \mathbf{v})\mathbf{v}$ denotes the centrifugal and Coriolis accelerations.

and the right-hand side is the resultant of applied forces:

- $\mathbf{g}(\mathbf{q})$ is the resultant of all conservative forces (only dependent on configuration). In this paper we consider gravity. Spring-like and magnetic forces can also be incorporated.
- $\mathbf{f}_{\text{ext}}(\mathbf{q})$ is the resultant of all generalized external non-conservative forces, excluding contact. An important example of this type of force are joint torques.




Figure 2-3: Distance to contact. For the example of a planar object falling on a horizontal surface, we consider the distances ϕ_n and ϕ_t between the closest point to the ground P and the origin O of a fixed reference frame. The point P is parametrized by the configuration of the object $\mathbf{q} = (x, y, \theta)$ and the angle β .

Using these equations, we may predict the trajectory (time-evolution) of the states for a prescribed set of external forces and initial conditions of the states. Rigid-bodies apply a force as they come into contact. We can consider the applied contact force as an external force to be added to the right hand side of the manipulator equations:

$$\mathbf{M}(\mathbf{q}) \frac{d\mathbf{v}}{dt} + \mathbf{C}(\mathbf{q}, \mathbf{v})\mathbf{v} = \mathbf{g}(\mathbf{q}) + \mathbf{f}_{\text{ext}}(\mathbf{q}) + \mathbf{J}_n^T(\mathbf{q})\mathbf{c}_n + \mathbf{J}_t^T(\mathbf{q})\mathbf{c}_t \quad (2.2)$$

With reference to Fig. 2-3, the new terms due to contact are:

- $\mathbf{J}_n^T(\mathbf{q}) = \nabla\phi_n(\mathbf{q})^T$, is the gradient of a "distance-to-contact" function $\phi_n(\mathbf{q})$ that determines the boundary between no contact ($\phi_n(\mathbf{q}) > 0$) and penetration ($\phi_n(\mathbf{q}) < 0$). It represents the contact normal vector in an inertial reference frame.
- $\mathbf{J}_n^T(\mathbf{q})\mathbf{c}_n$ is contact normal force in an inertial reference frame, where c_n is its magnitude.
- $\mathbf{J}_t^T(\mathbf{q})\mathbf{c}_t$ is the tangential contact force due to friction. c_t is the magnitude of the force along vectors in \mathbf{J}_t^T .

If we assume the contact forces are known, then eq. 2.2 is equivalent to eq. 2.1. The challenge is that the contact forces must be solved for along with the post con-

tact states. These contact forces must satisfy several physical constraints: i) non-penetration; ii) principal of maximum dissipation; and iii) conservation of energy.

The non-penetration constraint requires that the magnitude of the contact force be at least large enough such that the resultant relative normal velocity post contact is non-negative. The maximum dissipation principal requires that the frictional force oppose the relative object motion in the tangent plane. Conservation of energy requires that the net post contact energy does not exceed the net pre-contact energy of the system. Stewart and Trinkle [79] encoded these using complementarity constraints.

We may write the equations of motion and the constraints mathematically as (see [29] for detailed derivation):

$$\frac{d\mathbf{q}}{dt} = \mathbf{v}$$

$$\mathbf{M}(\mathbf{q})\frac{d\mathbf{v}}{dt} + \mathbf{C}(\mathbf{q}, \mathbf{v})\mathbf{v} = \mathbf{g}(\mathbf{q}) + \mathbf{f}_{\text{ext}}(\mathbf{q}) + \mathbf{J}_n^T(\mathbf{q})\mathbf{c}_n + \mathbf{J}_t^T(\mathbf{q})\mathbf{c}_t \quad (2.3)$$

$$\text{Subject to: } 0 = \mu\mathbf{J}_t(\mathbf{q})^T\mathbf{v} + \lambda\mathbf{e} \quad (2.4)$$

$$0 \leq \phi_n(\mathbf{q}) \perp \mathbf{c}_n \geq 0 \quad (2.5)$$

$$0 \leq \lambda \perp \mu\mathbf{c}_n - \sqrt{\mathbf{c}_{t,1}^2 + \mathbf{c}_{t,2}^2} \geq 0 \quad (2.6)$$

where $\mathbf{J}_t\mathbf{v}$ is the tangential velocity at the point of contact, $\mathbf{e} = \mathbb{1}_{n \times 1}$, and λ denotes the relative tangential velocity of the bodies at the point of contact.

The first constraint requires that the post contact velocity oppose the relative tangential velocity of the bodies at the point of contact. The second constraint requires that normal contact forces only exist when the distance function is zero. The last constraint requires that the tangential contact force is equal to the coefficient of friction multiplied by the normal contact force if the objects slide w.r.t to each other, else the tangential force must be less than this value (sticking contact). Here the tangential force is decomposed into two orthogonal vectors in the plane of contact.

The final constraint is conic (due to the friction cone). To simplify this constraint,

Stewart and Trinkle [79] used a polygonal approximation to the cone:

$$0 \leq \lambda \perp \mu \mathbf{c}_n - \sum_{i=1}^m c_{t,i} \geq 0 \quad (2.7)$$

where the net applied tangent force is the summation of the magnitude of the tangential forces along the positive span of the cone.

As a final step, the equations of motion and constraints are discretized in time using a first order semi-implicit Euler integration approximation (see [29, 79] for detailed derivation):

$$\begin{aligned} \mathbf{q}^{k+1} - \mathbf{q}^k &= h \mathbf{v}^{k+1} & (2.8) \\ \mathbf{M}(\mathbf{q}^{k+1}) (\mathbf{v}^{k+1} - \mathbf{v}^k) + \mathbf{C}(\mathbf{q}^{k+1}, \mathbf{v}^{k+1}) \mathbf{v}^{k+1} &= \mathbf{g}(\mathbf{q}^{k+1}) + \mathbf{f}_{\text{ext}}(\mathbf{q}^{k+1}) + \\ &+ \mathbf{J}_n^T(\mathbf{q}^{k+1}) \mathbf{c}_n^{k+1} + \mathbf{J}_t^T(\mathbf{q}^{k+1}) \mathbf{c}_t^{k+1} \\ 0 &\leq \phi_n(\mathbf{q}^k) \perp \mathbf{c}_n^k \geq 0 \\ 0 &\leq \lambda^k \perp \mu \mathbf{c}_n^k - (c_{t,1}^k + c_{t,2}^k) \geq 0 \\ 0 &\leq c_{t,1}^k \perp \lambda^k + v_{\text{rel},1}^k \geq 0 \\ 0 &\leq c_{t,2}^k \perp \lambda^k + v_{\text{rel},2}^k \geq 0 \end{aligned}$$

Here, the equations are written for the planar case, the same formulation holds for the 3D case. The first constraint constrains the normal contact force complement to contact distance. The second constraint complements the magnitude of the net tangent force (on or in the friction cone) to the relative sliding motion between objects. The last two constraints complement each component of the tangential force to the relative motion between objects, i.e. if the objects slip left, the tangential force point to the right is active and vice versa. Next, we will leverage this formulation to derive the contact-implicit system identification formulation.

2.3.2 Contact-implicit system identification

We present the derivation in four steps. Each step builds on the previous and handles one important aspect of the formulation:

Step 1 - Base inertial parameter set

The left hand side of the Manipulator equations (eq. 2.1) is a function of states and inertial parameters. We can re-written the left hand side as:

$$\mathbf{Y}(\mathbf{q}, \mathbf{v}, \dot{\mathbf{v}})\boldsymbol{\theta} = \mathbf{f}_{\text{ext}}(\mathbf{q}) \quad (2.9)$$

where $\mathbf{Y}\boldsymbol{\theta}$ is a refactoring of the left hand side into the product of two terms. The first term is a function of states and the second term, $\boldsymbol{\theta}$, is the *base inertial parameter* set. The base inertial parameters are the minimal identifiable set necessary to predict the dynamics of system given known external forces – for further details see [35, 47, 29]. One particularly useful application of the base inertial parameter set in robotics is adaptive control of manipulators [76]. In our derivation, we re-parameterize the equations of motion in the time-stepping formulation in eq. 2.8 to reflect the base inertial parameters:

$$\begin{aligned} \mathbf{q}^{k+1} - \mathbf{q}^k &= h\mathbf{v}^{k+1} & (2.10) \\ \mathbf{Y}(\mathbf{q}^{k+1}, \mathbf{v}^k, \mathbf{v}^{k+1})\boldsymbol{\theta} &= \mathbf{f}_{\text{ext}}(\mathbf{q}^{k+1}) + \mathbf{J}_n^T(\mathbf{q}^{k+1})\mathbf{c}_n^{k+1} + \mathbf{J}_t^T(\mathbf{q}^{k+1})\mathbf{c}_t^{k+1} \\ 0 &\leq \phi_n(\mathbf{q}^k) \perp \mathbf{c}_n^k \geq 0 \\ 0 &\leq \lambda^k \perp \mu\mathbf{c}_n^k - (c_{t,1}^k + c_{t,2}^k) \geq 0 \\ 0 &\leq c_{t,1}^k \perp \lambda^k + v_{\text{rel},1}^k \geq 0 \\ 0 &\leq c_{t,2}^k \perp \lambda^k + v_{\text{rel},2}^k \geq 0 \end{aligned}$$

Note that the constraints are left unchanged. This is because only the dynamic equations of motion are directly a function of inertia, the contact forces are decided as a function of these parameters during optimization.

Step 2 - Contact forces

Next, we handle the contact forces. Considering eq. 2.10, we note that the Jacobians projecting the contact forces to the generalized coordinates are a function of states, not parameters or forces. As such we re-write eq. 2.10:

$$\begin{aligned}
 \mathbf{q}^{k+1} - \mathbf{q}^k &= h\mathbf{v}^{k+1} & (2.11) \\
 \left[\mathbf{Y}(\mathbf{q}^{k+1}, \mathbf{v}^k, \mathbf{v}^{k+1}) \quad -\mathbf{J}(\mathbf{q}^{k+1}) \right] \begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{c}^{k+1} \end{bmatrix} &= \mathbf{f}_{\text{ext}} \\
 0 \leq \phi_n(\mathbf{q}^k) \perp \mathbf{c}_n^k &\geq 0 \\
 0 \leq \lambda^k \perp \mu \mathbf{c}_n^k - (c_{t,1}^k + c_{t,2}^k) &\geq 0 \\
 0 \leq c_{t,1}^k \perp \lambda^k + v_{\text{rel},1}^k &\geq 0 \\
 0 \leq c_{t,2}^k \perp \lambda^k + v_{\text{rel},2}^k &\geq 0
 \end{aligned}$$

where the left hand side is now factored into the product of two terms. The first term is only a function of states ($\mathbf{A}^k = \left[\mathbf{Y}(\mathbf{q}^{k+1}, \mathbf{v}^k, \mathbf{v}^{k+1}) \quad -\mathbf{J}(\mathbf{q}^{k+1}) \right]$) and the second is the concatenation of the unknown base inertial parameters and contact forces.

Step 3 - A window of observations

The dynamic equations of motion and constraints in eq. 2.11 are written for time instant t . To effectively reason over the set of parameters, we need to consider a window of observations and compute the most likely estimates of these values such that the observations in this window are best explained.

To this end, we may write eq. 2.11 over a window of observations as the optimiza-

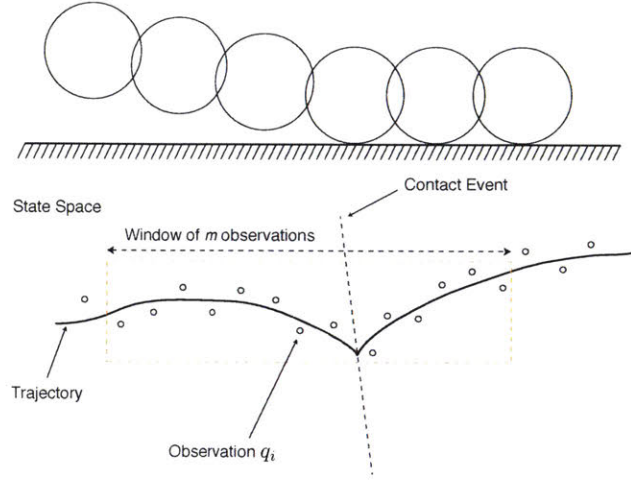


Figure 2-4: A disk transitions from free flight to rolling contact. The solid line denotes the true trajectory of the disk, and the dots denote observations corrupted by sensor noise. Given an observation window of length m we want to estimate Θ and $\{c^1 \dots c^m\}$ for the duration of the window.

tion program, Fig. 2-4:

$$\begin{aligned}
 & \underset{\theta, c_{1..m}}{\text{minimize}} && \left\| \mathbf{f}_{\text{ext}}^{k+1} - \mathbf{A}^k \begin{bmatrix} \theta \\ \mathbf{c}^{k+1} \end{bmatrix} \right\|_2 && (2.12) \\
 & \text{s.t.} && 0 \leq \phi_n(\mathbf{q}^k) \perp \mathbf{c}_n^k \geq 0 \quad \forall k = 1..m \\
 & && 0 \leq \lambda^k \perp \mu \mathbf{c}_n^k - (c_1^k + c_2^k) \geq 0 \quad \forall k = 1..m \\
 & && 0 \leq c_1^k \perp \lambda^k + v_{\text{rel},1}^k \geq 0 \quad \forall k = 1..m \\
 & && 0 \leq c_2^k \perp \lambda^k + v_{\text{rel},2}^k \geq 0 \quad \forall k = 1..m
 \end{aligned}$$

Intuitively, this optimization program minimizes the error in dynamics by computing the best estimate of the parameters and contact forces to match the sequence of observations. The cost function is a convex quadratic expression. The constraints are linear complementarity.

The \mathbf{A} regressor is composed of the concatenation of \mathbf{Y} with the contact Jacobian at each time-step. The parameter vector θ has been appended with the contact forces for each time-step. The constraints resolve the contact mode and enforce physically realizable values for the contact forces as well as the inertial parameters.

If \mathbf{f}_{ext} is sufficiently rich, and the system is in contact, then the identifiable set is the full vector $\boldsymbol{\theta}$ as well as all contact forces. The state dependant contact constraints do not affect the base inertial parameter set, they instead impose that the estimated contact forces and identified parameters are physically consistent.

Step 4 - Handling uncertainty for LCP constraints

As written, the optimization program in eq. 2.12 has an important short-coming. When discretizing and/or when dealing with noisy measurements, the probability of sampling on the contact manifold is zero. The distance to contact is never exactly zero, therefore the normal contact force cannot take on a non-zero value. Consequently, the estimation yields unreasonable parameter values to try to explain contact observations without employing contact forces.

To address this issue, we may present three solutions:

1. Soften the distance constraint,
2. Reformulate to minimize complementarity error,
3. pre-process data to using projections to the contact manifold.

Approach 1 - Soften the distance constraint

In this approach, rather than requiring the distance complementarity constraint to be satisfied precisely, we penalize its violations in the cost:

$$\begin{aligned}
 & \underset{\boldsymbol{\theta}, \mathbf{c}_{1..m}}{\text{minimize}} \quad \left\| \mathbf{f}_{\text{ext}}^{k+1} - \mathbf{A}^k \begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{c}^{k+1} \end{bmatrix} \right\|_2 + \mathbf{w}^T (\Phi_n \cdot \mathbf{C}_n) & (2.13) \\
 & \text{s.t:} \quad \phi_n(\mathbf{q}^k) \geq 0, \quad \mathbf{c}_n^k \geq 0 \quad \forall k = 1..m \\
 & \quad 0 \leq \lambda^k \perp \mu \mathbf{c}_n^k - (\mathbf{c}_1^k + \mathbf{c}_2^k) \geq 0 \quad \forall k = 1..m \\
 & \quad 0 \leq \mathbf{c}_1^k \perp \lambda^k + v_{\text{rel},1}^k \geq 0 \quad \forall k = 1..m \\
 & \quad 0 \leq \mathbf{c}_2^k \perp \lambda^k + v_{\text{rel},2}^k \geq 0 \quad \forall k = 1..m
 \end{aligned}$$

Intuitively, the optimization now allows for forces at a small distance from the contact. The amount is penalized by the weight vector \mathbf{w} . The cost function is still quadratic and convex. We maintain the constraint that distance and normal contact forces must be non-negative.

Approach 2 - Complementarity constraints as costs

In this approach, we formulate the complementarity constraints as the cost and require that the dynamics hold as a constraint instead. This further relaxes the strict complementarity conditions of sticking and sliding to account for some tangential uncertainty. We may write:

$$\begin{aligned}
& \underset{\boldsymbol{\theta}, \mathbf{c}_{1..m}}{\text{minimize}} && \mathbf{w}_1^T (\Phi_n \cdot \mathbf{C}_n) + \mathbf{w}_2^T (\boldsymbol{\lambda} \cdot \mathbf{F}_c) + \mathbf{w}_3^T (\mathbf{C}_t \cdot \mathbf{V}_r) && (2.14) \\
& \text{s.t:} && \mathbf{f}_{\text{ext}}^{k+1} - \mathbf{A}^k \begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{c}^{k+1} \end{bmatrix} = 0 \quad \forall k = 1..m \\
& && \phi_n(\mathbf{q}^k) \geq 0, \quad \mathbf{c}_n^k \geq 0 \quad \forall k = 1..m \\
& && 0 \leq \lambda^k \perp \mu \mathbf{c}_n^k - (c_1^k + c_2^k) \geq 0 \quad \forall k = 1..m \\
& && 0 \leq c_1^k \perp \lambda^k + v_{\text{rel},1}^k \geq 0 \quad \forall k = 1..m \\
& && 0 \leq c_2^k \perp \lambda^k + v_{\text{rel},2}^k \geq 0 \quad \forall k = 1..m
\end{aligned}$$

where the matrices \mathbf{C}_n , \mathbf{F}_c , \mathbf{C}_t , and \mathbf{V}_r denote the concatenation of the normal contact forces, friction cone, tangential friction force, and relative sliding vectors, respectively. Note that we may also soften the dynamics by placing them in the cost function and penalizing dynamics violations.

Approach 3 - Contact manifold projections

In this approach, rather than requiring the optimization to handle to uncertainty, we pre-filter the data by projecting the noisy and discretized trajectory to the closest contact manifold. The projection then allows the constraints to be precisely satisfied and leads to a numerically well-posed optimization.

Fig. 2-5 illustrates how to implement this approach. An important challenge

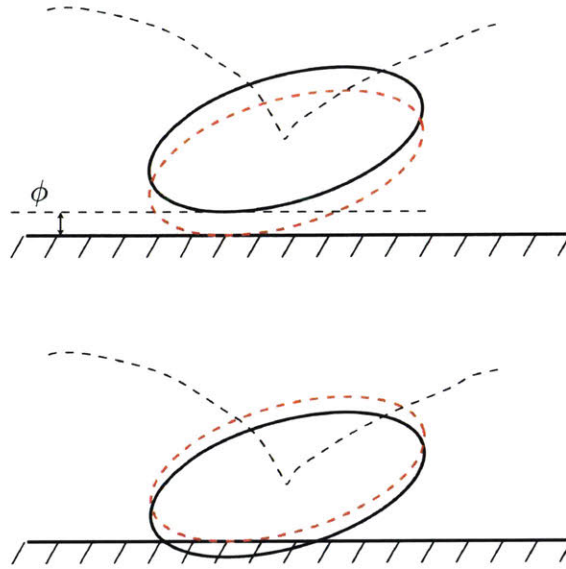


Figure 2-5: Uncertainty in the pose measurement can lead to configurations that do not lie on the contact manifold. Projection algorithms or relaxation in the constraints can be used to account for this uncertainty

in implementing this filtering technique is effective contact detection such that the correspondence in the projection matches the contact events accurately.

2.3.3 Identifiability analysis

In its barest form system identification seeks to answer two questions, given a dynamic model of a phenomena and data collected from the responses of the system to a given excitation, i) what are the set of the parameters from the dynamic model can we can hope to estimate (identifiability), and ii) how to estimate them (identification).

Classical identifiability analysis starts from the observation that the left-hand side of the dynamic equation of motion in Eq. 2.10 can be written as a linear map $\mathbf{Y}\boldsymbol{\theta}$ of the base-inertial parameter vector. The structure of the base inertial parameters depends on the geometry of the problem, and dictates what can or cannot be identified or estimated. We will present an example shortly.

It is important to note that the identification is only effective if the external forces are informative enough to excite the complete spectrum of dynamics of the system considered, i.e. persistence of excitation criteria for the actuation signal is met:

$$\alpha \mathbf{I} \leq \frac{1}{T} \int_t^{t+T} \mathbf{Y}(\tau)^T \mathbf{Y}(\tau) d\tau$$

for the observed trajectories to be informative of the parameters to identify. The parameter α is a chosen positive scalar and I denotes the identity matrix. Intuitively this criteria quantifies how positive definite the matrix $\mathbf{Y}^T \mathbf{Y}$ is over the window of observation of length T . Loosely put, the importance of the positive definiteness of this matrix is due to the inversion it undergoes in estimating the parameters over the time period. If \mathbf{f}_{ext} meets this criteria and the system is in contact, then the identifiable set is the full vector $\boldsymbol{\theta}$ as well as all contact forces.

Regarding the ID optimization program (Eq. 2.12) the persistence of excitation criteria sets conditions on the \mathbf{Y} and \mathbf{J} matrices respectively. To illustrate, we write:

$$\begin{aligned} \left[\mathbf{Y}(\mathbf{q}^{k+1}, \mathbf{v}^k, \mathbf{v}^{k+1}) - \mathbf{J}(\mathbf{q}^{k+1})^T \right] \begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{c}^k \end{bmatrix} &= \mathbf{f}_{\text{ext}}^{k+1} \\ \begin{bmatrix} \boldsymbol{\theta} \\ \mathbf{c}^k \end{bmatrix} &= \left[\mathbf{Y}(\mathbf{q}^{k+1}, \mathbf{v}^k, \mathbf{v}^{k+1}) - \mathbf{J}(\mathbf{q}^{k+1})^T \right]^\dagger \mathbf{f}_{\text{ext}}^{k+1} \end{aligned}$$

The pseudo inverse of the regressor matrix is numerically well-conditioned when the criteria for persistent excitation of the signals is satisfied. Fig. 2-5 shows an intuitive example, the contact of a square against a flat surface. In the left image the square lands perfectly on one side whereas on the right it lands at an angle. The left square will continue to bounce up and down vertically without ever rotating, implying no information can be deduced regarding its rotational inertia (in fact it cannot be distinguished from a point particle). The square on the right rotates with various angular velocities as it bounces, which provides information regarding its rotational inertia, so the excitation can be considered rich.

In the case $\mathbf{f}_{\text{ext}} = \mathbf{0}$, i.e., there is no known external actuation on the system for the duration of the motion. In this case the system in Eq. 2.12 becomes homogeneous, and the optimization reduces to a singular value problem. The contact forces cannot

be found independently of the parameter vector $\boldsymbol{\theta}$ and only a ratio of the two can be found.

The gravitational force does not yield additional information with respect to the identifiability of inertial parameters. This is due to the fact that these terms are linear in masses and are consequently proportional to accelerations. An intuitive example is that of a free falling object where $m\mathbf{g} = m\mathbf{a}$, and observing its motion is not informative about its mass, a notion that extends to articulated rigid bodies.

When there is no contact, i.e., $\phi_n(\mathbf{q}) > 0$, all constraints are satisfied trivially, and the nonlinear optimization problem becomes again linear least squares regression over the parameter vector $\boldsymbol{\theta}$. This is the case for several seminal works on system identification applied to robotics [47, 76].

The optimization program in Eq. 2.12 is formulated to deal with unknown contact modes, this is critical since in practice these modes are often not known beforehand. Here, we demonstrate identifiability on a single rigid body making contact with a rigid flat surface and solve away the complementarity constraints for the cases of sliding contact. See [29] for the sticking case, both cases result in the same set of identifiable parameters.

Here we assume that the system of interest is a single rigid body making contact with a rigid and flat surface and resolve the complementarity constraints by assuming it is in sliding contact mode during the window of observation. This breaks the hybridness that characterizes the dynamics of intermittent frictional contact. During sliding the complementarity constraints in Eq. 2.12 become:

$$\mathbf{c}_n^k > 0 \quad \lambda^k > 0 \quad \mathbf{c}_t^k > 0 \quad (2.15)$$

where \mathbf{c}_t^k is $nvecc_1^k$ or $nvecc_2^k$, depending on whether the object is sliding to the right or the left.

The choice of contact mode turns the problem into an unconstrained optimization. Expanding Eq. 2.12 with the general expressions in Eq. 2.8, and rewriting the linear system with \mathbf{c}_n , \mathbf{c}_t and λ as variables, we obtain:

$$\begin{bmatrix} \mathbf{J}_n^T \mathbf{M}^{-1} \mathbf{J}_n & \mathbf{J}_n^T \mathbf{M}^{-1} \mathbf{J}_t & 0 \\ \mathbf{J}_t^T \mathbf{M}^{-1} \mathbf{J}_n & \mathbf{J}_t^T \mathbf{M}^{-1} \mathbf{J}_t & 1 \\ \mu & -1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{c}_n^k \\ \mathbf{c}_t^k \\ \lambda^k \end{bmatrix} = - \begin{bmatrix} \mathbf{J}_n^T \mathbf{b} \\ \mathbf{J}_t^T \mathbf{b} \\ 0 \end{bmatrix} \quad (2.16)$$

where:

$$\mathbf{b} = \mathbf{v}^k + h\mathbf{M}^{-1}(-\mathbf{g}(\mathbf{q}^k) - \mathbf{C}(\mathbf{q}^k, \mathbf{v}^k) + \mathbf{f}_{\text{ext}}^k)$$

Solving the linear system of equations for \mathbf{c}_n and \mathbf{c}_t we can write:

$$\mathbf{c}_n^k = m \frac{v_y^k + \theta^k J_y^k - hg + \frac{h}{m}(f_y^k + h\frac{m}{I}f_\theta)}{1 + \left(J_y^{k2} + \mu J_x^k J_y^k\right) \frac{m}{I}} = \frac{\mathbf{c}_t^k}{\mu} \quad (2.17)$$

which solves contact forces as functions of the inertial properties, geometry of contact, and kinematic measurements. We replace the expressions for normal and tangential contact forces back into the equations of motion, yielding:

$$\begin{bmatrix} v_x^{k+1} - v_x^k \\ v_y^{k+1} - v_y^k + hg \\ v_\theta^{k+1} - v_\theta^k \end{bmatrix} = \frac{v_y^k + v_\theta^k J_y^k - hg + \frac{h}{m}(F_y^k + h\frac{m}{I}F_\theta^k)}{1 + \left(J_y^{k2} + \mu J_x^k J_y^k\right) \frac{m}{I}} \begin{bmatrix} \mu \\ 1 \\ \frac{m}{I}(J_y^k + \mu J_x^k) \end{bmatrix} + h \begin{bmatrix} \frac{F_x^k}{m} \\ \frac{F_y^k}{m} \\ \frac{F_\theta^k}{I} \end{bmatrix} \quad (2.18)$$

In order to study the identifiability of the system making contact we need to manipulate this expression into a linear form in the unknown parameters. With some algebraic manipulation, we can write $\mathbf{Y} \cdot \boldsymbol{\theta} = \mathbf{f}$, with:

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}_1(\mathbf{q}, \mathbf{v}) & \mathbf{y}_2(\mathbf{q}, \mathbf{v}, \mathbf{f}_{\text{ext}}) & \mathbf{y}_3(\mathbf{q}, \mathbf{v}, \mathbf{f}_{\text{ext}}) & \mathbf{y}_4(\mathbf{q}, \mathbf{v}, \mathbf{f}_{\text{ext}}) \end{bmatrix}$$

$$\boldsymbol{\theta} = \begin{bmatrix} \frac{m}{I} & \frac{1}{m} & \frac{1}{I} & \frac{m}{I^2} \end{bmatrix}^T$$

where \mathbf{y}_i denote the collections of the coefficients of the system parameters. This linear mapping $\mathbf{f} = \mathbf{Y}\boldsymbol{\Theta}$ in the base parameters $\boldsymbol{\Theta} = \left[\frac{1}{m} \frac{1}{I} \frac{m}{I} \frac{m}{I^2}\right]^T$ indicates that

this is indeed the identifiable set, i.e. that provided we have sufficiently rich external excitation, we can estimate m and I independently. To estimate these parameters we collect N data points of the states and external forces and solve a linear least squares optimization to evaluate the mass and inertia of the object then plug the values found back into the previous expression to find the contact forces.

A key observation is that \mathbf{Y} is dependant on the states as well as the external forces, but the second, third and fourth column are linear functions of the external forces whereas the first column is only a function of the states. This implies that if there is no known external excitation during the contact then $\mathbf{y}_2 = \mathbf{y}_3 = \mathbf{y}_4 = \mathbf{0}$ and only the first column of the \mathbf{Y} matrix provides any information, consequently only the first element of $\boldsymbol{\theta}$ can be estimated, i.e. the ratio of the mass to the inertia. We can then compute the ratio of the contact forces to mass using this value.

So far we have not studied any particular contact geometry, and assumed that the contact Jacobian is a sufficient representation of the body. Yet algebraic expressions are required to compute that Jacobian as a function of the configuration of the object and its geometry. The following subsections show examples of analytical expressions for a block, ellipse, and rimless wheel which we will later use in simulated and real experiments.

2.4 Experiments - Simulation

In this section, we first derive contact Jacobians and the optimization for system identification for three planar geometric shapes. Next, we present results from simulations of planar single and multi-body frictional interactions.

2D Block

The block in Fig. 2-6 models as a square with side length a , orientation θ , center of mass (x, y) , mass m , and angular inertia I . The minimum distance from the block to




Figure 2-6: Block in free-fall.

the ground is the minimum of the vertical distances from all vertex:

$$\phi_n = \min \left(\begin{array}{c} y - \frac{a}{\sqrt{2}} \cos(\pi/4 - \theta) \\ y - \frac{a}{\sqrt{2}} \cos(\pi/4 + \theta) \\ y + \frac{a}{\sqrt{2}} \cos(\pi/4 - \theta) \\ y + \frac{a}{\sqrt{2}} \cos(\pi/4 + \theta) \end{array} \right) \quad (2.19)$$

The contact Jacobian is derived from differentiating the distances in eq. 2.19. Assuming the first vertex is the lowest (with $\beta = \pi/4$ and $l = a/\sqrt{2}$), and since curvature does not play a role, i.e., $\frac{\partial \beta}{\partial \theta} = 0$, the expression of the normal and tangential Jacobians become:

$$\mathbf{J}_n(\mathbf{q}) = \left[0 \quad 1 \quad -\frac{a}{\sqrt{2}} \sin(\pi/4 - \theta) \right], \quad \mathbf{J}_t(\mathbf{q}) = \left[1 \quad 0 \quad \frac{a}{\sqrt{2}} \cos(\pi/4 - \theta) \right] \quad (2.20)$$

which we could have also obtained by directly differentiating the distance function.

2D Ellipse

The ellipse in Fig. 2-7 incorporates extra complexity due to its curvature. We will see that the resulting contact dynamics are very sensitive to orientation, because small perturbations can produce large changes in contact location.




Figure 2-7: Ellipse in free fall.

We parametrize the perimeter curve of the ellipse by angle β . Denote the major and minor radii of the ellipse with a and b , and its orientation with respect to a fixed reference frame by θ . The relationship between β to θ is given by:

$$\tan(\pi - \theta) = -\frac{b}{a} \cot \beta \xrightarrow{\frac{\partial}{\partial \theta}} \frac{\partial \beta}{\partial \theta} = -\frac{a}{b} \frac{1 + \tan^2(\pi - \theta)}{1 + \cot^2(\beta)} \quad (2.21)$$

We now write the distance of any point on the perimeter of an ellipse from its center as:

$$l(\beta) = \frac{ab}{\sqrt{b^2 \cos^2 \beta + a^2 \sin^2 \beta}} \xrightarrow{\frac{\partial l}{\partial \theta}} \frac{\partial l}{\partial \theta} = \frac{ab(b^2 - a^2) \sin 2\beta}{2\sqrt{(b^2 \cos^2 \beta + a^2 \sin^2 \beta)^3}} \frac{\partial \beta}{\partial \theta} \quad (2.22)$$

Which allows us to find the expression for $\partial l / \partial \theta$. By back-substituting we can find an expression for the Jacobian.

Rimless Wheel

The rimless wheel is a classical example of a very simple passive walker. Its stability and gait cycles have been extensively studied. Here we demonstrate the application of parameter and contact force estimation approach to a rimless wheel as it “walks” down an incline. We make no assumption of sticking vs sliding motion. We may




Figure 2-8: Rimless wheel rolling/sliding downhill

write:

$$\min_{\theta, c_1 \dots c_m} \left\| \begin{bmatrix} v_x^{k+1} - v_x^k \\ v_y^{k+1} - v_y^k \\ v_\theta^{k+1} - v_\theta^k \end{bmatrix} - h \begin{bmatrix} \frac{1}{m} & 0 & 0 \\ 0 & \frac{1}{m} & 0 \\ 0 & 0 & \frac{1}{I} \end{bmatrix} \left(\begin{bmatrix} 0 \\ -mg \\ 0 \end{bmatrix} + \sum_{i=1}^p c_{n,i}^{k+1} \mathbf{J}_{n,i}^{k+1} + c_{t,i}^{k+1} \mathbf{J}_{t,i}^{k+1} \right) \right\|_2 \quad (2.23)$$

where we are now summing over the contact forces at all legs/spokes, since all legs can make contact with the ground. Note also the increase in complexity due to the fact that two simultaneous contact forces are possible.

Results from Simulated Data

To demonstrate the identification procedure we start from simulated data originated by a time-stepping LCP simulator based on the linear formulation in Eq. 2.3. In the simulation we use a unit mass (kg), coefficient of restitution of 0.6 and a horizontal flat rigid surface with coefficient of friction 0.7. The ratio of mass m to angular inertia I moment of inertia for the block and ellipse are 6 (m^2) and 0.8 (m^2) respectively.

To generate data we gave the bodies a random set of initial positions and velocities and simulate their trajectory 100 times. Fig. 2-9a, Fig. 2-9b, and Fig. 2-9c show traces of example trajectories for the block, ellipse, and rimless wheel respectively. Fig. 2-10 illustrate the dependence of the motion of the rimless wheel with its mass.

Before testing the estimation algorithm we add Gaussian noise $\sim \mathcal{N}(0, \sigma^2)$ to the

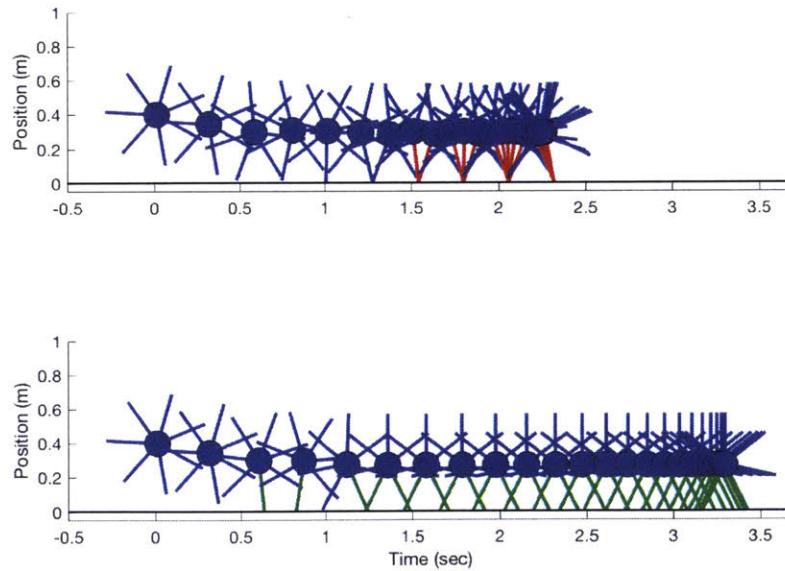
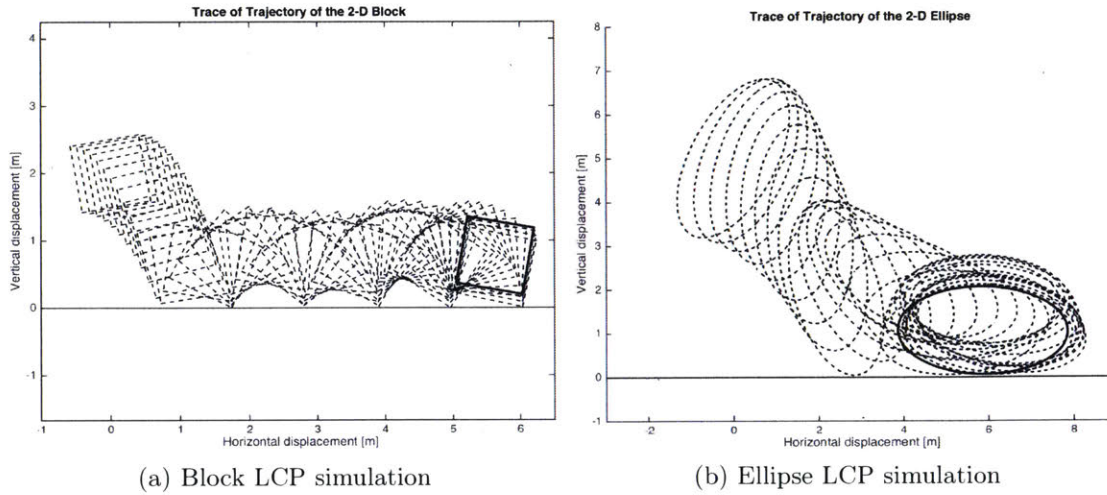


Figure 2-9: a) LCP simulation trace of a falling block. b) LCP simulation trace of a falling ellipse. c) LCP simulation trace of a falling rolling rimless wheel for the cases of high friction (top) and low friction (bottom) with the ground.




Figure 2-10: Dependence of the motion of the rimless wheel with its mass.

Table 2.1: 2D Numerical Simulation and Identification Results

Block		Ellipse		Noise
Mean Error (m^2) \pm Std.	% Error	Mean Error (m^2) \pm Std.	% Error	S.N.R. (dB)
-0.021 ± 0.097	-0.35	-0.031 ± 0.101	-0.31	40
-0.125 ± 0.311	-2.083	-0.136 ± 0.443	-2.34	30
-0.659 ± 0.912	-10.98	-0.712 ± 0.820	-11.71	20
-3.491 ± 1.366	-58.197	-4.891 ± 1.938	-80.44	10

simulated data (both configurations q and velocities v) where σ is a function of signal to noise ratio. We then use the resulted signals to estimate $\frac{m}{I}$.

Tab. 2.1 shows the results of the optimization as a function of signal to noise ratio where the first column denotes the mean error between predicted and actual $\frac{m}{I}$ value, the second column denotes the percent error and the final column denotes the magnitude of the noise on measurements. We see good agreement between the predicted and true parameter with low levels of noise and a steady deterioration of prediction as noise is increased. We attribute the increasing error mostly to the effect that configuration has in the determination of the contact Jacobian. Poor evaluation of these variables results in poor behavior prediction, which makes it difficult to

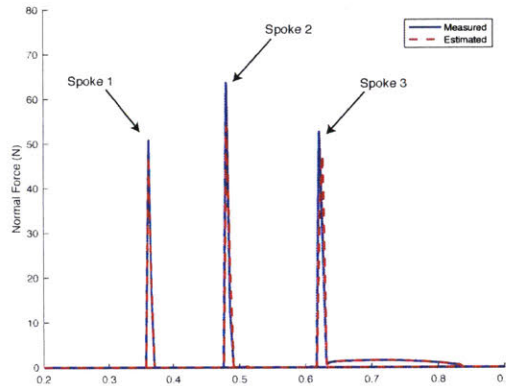


Figure 2-11: Normal contact force of the rimless wheel spokes as a function of time.

estimate parameters.

For the rimless wheel system we show the estimated contact force to demonstrate the algorithm's ability to recover these forces assuming known mass. Fig. 2-11 shows the estimation of the contact force for the case where the rimless wheel rolls without sliding coming from the first 3 spokes, as can be seen the algorithm can recover the reaction force closely.

Identification for Multi-Body Interactions

In this section we demonstrate the application of the technique to a multi-body, multi-contact problem. Fig. 2-1 shows a two link manipulator pushing a disk resting on flat ground. The disk is in contact both with the manipulator and the environment. Initially the manipulator and the disk are not in contact and the disk starts to drop from a height, subsequently the two link manipulator makes contact with the disk and pushes the disk against the ground, the disk begins to roll and slide its way out of the wedge formed by the manipulator and ground.

The system has a total of five degrees of freedom (two rotational joints of the manipulator and the position and orientation of the disk). Jointly with their velocities, it makes for a 10-dimensional state space. We assume a maximum of 2 possible contact pairs at any given instant, the manipulator-ball pair and the ball-ground pair, and for

each pair there are the associated normal and tangential contact forces in the contact frame, as well as their corresponding complementarity constraints. The control inputs are the joint torques of the manipulator, and everything is under the influence of a vertical gravitational field.

Assuming a window length of m , the total number of parameters to estimate is $4m + 3 + 2$ where there are $2m$ forces per contact pair, 3 base inertial parameters of the two-link manipulator and 2 base inertial parameters of the ball. Note that the number of parameters to optimize grows linearly with the length of the window and there will be an additional $2m$ per any new contact pair.

To simulate the experiment we initialize the manipulator and disk in a configuration similar to Fig. 2-1 and apply pre-determined torque profiles to the two-link manipulator. We log the resulting states and contact forces generated. Next we use the recorded configurations and torque profiles as inputs to the optimization along with the complementarity constraints to estimate the contact forces and base inertial parameters. Fig. 2-12 shows estimated contact forces in the three directions between the disk and the manipulator overlaid with the true values.

2.5 Experiments - Real World

To further validate the approach we constructed the experimental setup in Fig. 2-13. This setup is designed to emulate planar physical interaction closely and match the simulations.

The dropping arena is constructed with two very flat sheets of glass with support spacers to constraint the motion of a falling object in the plane. The objects are 3D printed in hard plastic with a square of side length 70 mms and ellipse with major and minor axes lengths of 70 and 50 mms. We track the position of the objects with an accurate Vicon motion tracking system at 250 Hz which proved accurate enough to extract velocity estimates by low-pass filtering and differentiation. To collect ground truth measurements of force we used an ATI Gamma 6 axis force torque sensor with 1000 Hz sampling rate. For each drop experiment we considered the first 3 bounces,




Figure 2-12: Overlay of estimated and measured contact force between the manipulator and disk. The contact forces and inertial parameters are estimated concurrently. The contact forces are estimated to within approximately 20%. The discrepancy is partially attributed to the short time horizon of the interaction providing limited persistence of excitation.




Figure 2-13: Experimental dropping setup. A six axis industrial robot latches magnetically on a part, gives it some initial velocity and orientation and drops it. A motion capture system and a F/T sensor capture the falling motion at high frequency.




Figure 2-14: Example of a falling trajectory of a planar ellipse in the experimental dropping arena. The figure shows a very sparse sub-sample of the trajectory recorded by the motion tracking system.

which are extracted automatically using the impact signature captured by the F/T sensor. Fig. 2-14 shows an example of such a trajectory.

We used the data from a total of 280 drop tests to evaluate the proposed algorithm to estimate the ratio of mass to inertia of the object, as well as the peak contact forces at contact. The ratio of mass to inertia (ground truth $535.03 m^2$) was estimated as $587.50 m^2$ (standard deviation in error of 27.6) and mean error in peak contact force was -12.12 Newton (standard deviation in error of 31.59) where on average the peak contact force was approximately 135 Newtons.

Fig. 2-15 shows an example of contact force estimation. The ellipse is shown as it comes into contact with the surface and the reaction force that is predicted by the algorithm overlaid with the ground truth. Note that the contact force, as measured by the F/T sensor, shows oscillations post impact. Our hypothesis is that these oscillations are artifacts due to oscillations of the impact plate and the internal compliance of the force/torque sensor (strain gauges). Though secondary impacts are a possible cause of the oscillations, we could not detect their occurrence from the state observations – most contact events appear clean single impacts from the state observations. For these reason we focus on recovering the contact forces at impact.

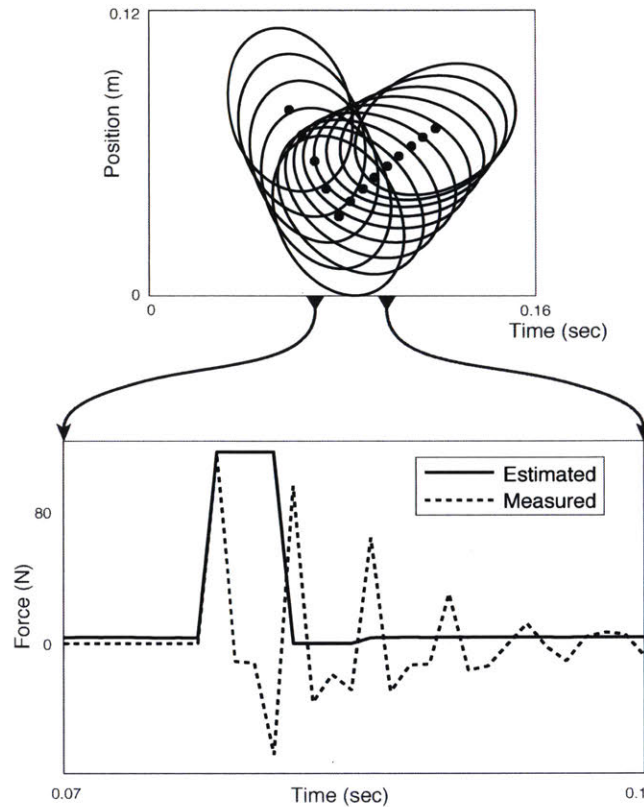


Figure 2-15: Reconstruction of the contact force for a impact of an ellipse.

Part of the error in peak force estimation may be due to these unmodeled dynamics.

2.6 Related Work

System identification studies the problem of fitting a model (i.e., inertial parameters) to a series of inputs (i.e., forces and torques) and responses (i.e., displacements/velocities/accelerations) of a dynamic system. The basic idea behind system identification is that, although the response of a dynamic system tends to be complex, the governing dynamics are often linear in a set of observable parameters. For example, while $\sum \mathbf{f} = m \cdot \mathbf{a}$ can lead to complex trajectories, forces and accelerations are still linearly related by m . In unconstrained dynamic systems, this allows closed-form least-squares formulations for the estimation of those parameters.

System identification determines what parameters are instrumental to a particular

dynamic system, what observations are informative, and how to excite the system to trigger those observations, ultimately yielding estimates of the parameters from measured data. This idea has been applied in robotics to the identification of serial and parallel link manipulators [35, 47], and to identify inertial parameters sufficient for control purposes [76].

In robotic manipulation, we often rely on dynamic models of impact and frictional interaction by assuming known masses, inertias, and friction or restitution coefficients, among others. For example many algorithms for contact-aware state estimation ([26], [6], [52], [94], [92]), use a dynamic model to filter noisy observations of state. [93] studies the problem of simultaneously estimating state and inertial and frictional parameters in a planar pushing task. Our work provides the basis to understand what parameters can we hope to estimate from the type of interaction and the available information. [7] exploit the particularities of floating base robots to simplify the estimation of their inertial parameters. In this work we set the focus on dealing with the issues that originate from the hybridness of making and breaking contact without making assumptions about the particular platform.

Closest to our work, [51] develop a similar approach to estimating inertial parameters, based on a physics engine with smoothed-out contact dynamics, leading to computationally efficient algorithms. More recently, in the context of planar pushing task, [95] propose to constrain the search for dynamic models to a convex polynomial relationship between forces and motions. The assumption, motivated by the principle of maximal dissipation and the concept of limit surface ([39]), yields a data-efficient algorithm. In a similar spirit, we are interested in studying under what circumstances those parameters are identifiable as a function of the available sensor data and the type of interaction.

Algorithms for planning and control of dynamic manipulation through contact ([57], [69], [71], [19]) also rely heavily on known dynamic parameters to make prediction for given control inputs.

Dynamic models are also widely used for fault detection and task monitoring, although to a lesser degree within the context of frictional contact. [90] provides an

early review of methods to detect and diagnose changes in the evolution of a system from its expected behavior, where the expected behavior can be specified as following a particular parametrized dynamic model ([8], [74],[21]) or data-driven model ([83], [73]). Recently, [58] propose to use a particle filter based scheme to estimate and localize external contacts to explain unexpected internal torques in a floating base mechanism like a humanoid. In all these cases, system identification has the potential to provide a formal approach to extract estimates of system parameters and contact forces from observations of rigid-body contact interactions.

The main design choice in our approach is the selection of a time-stepping Linear Complementarity Problem (LCP) scheme for the resolution of forces and accelerations (or rather impulses and velocities) during frictional contact. Why LCP? [14] identifies 3 classes of methods for rigid body simulation:

- *Penalty methods* approximate contact interactions by allowing interpenetration of bodies and generating forces proportional to the amount of penetration, yielding smooth yet stiff systems that can be solved with integration methods. These methods yield computationally favorable solutions at the cost of realism.
- *Event-driven methods* rely on a listing, resolution, and selection of all possible contact/impact events. They typically require some knowledge of contact time which may be difficult to predict *a priori* for complicated multi-body systems.
- *Time-stepping methods* integrate the equations of motion during a finite time interval. Should a contact (or multiple) be detected during the interval, the algorithm resolves the collisions and continues to integrate the equations of motion.

The time-stepping approach, in conjunction with the velocity-impulse resolution of contact, which results in a Complementarity Problem (CP), has been advocated by [80] and [5] among others, and has been shown to be robust to phenomena such as Painleve’s problem [78], and to always yield a solution, with linear approximations of the friction cone. Specially interesting for this work, LCP, or CP in general, provides

a unique consistent formulation over different contact modalities, i.e., contact vs. separation and sticking vs. slipping.

2.7 Discussion

In this chapter, we studied the estimation of contact forces and inertial parameters for systems of rigid bodies undergoing planar frictional contact. We formulated the estimation as a nonlinear optimization problem where the various contact modes are handled using complementarity constraints. We also discussed the identifiability and identification of parameters and forces for systems with and without the presence of external forces. This approach provides one method of uniformly handling the hybrid modes due to contact under the assumptions of rigid-body contact constraints.

We expect this approach to work fairly well when the system under study well approximated by the rigidity assumptions. Predictions using contact free rigid-body motion models (manipulator equations) are fairly accurate. The contact models may introduce some errors in the estimations made by the approach. These errors may be due to approximations made in the contact models for computational-efficiency. Deviations from the assumptions of the models such as significant deformation during contact and finite-time duration of contact events may contribute to the quality of the estimations. Some of the empirical results we found in this chapter suggest that some of these factors may be at play. We dedicate the next chapter to investigating the limitations of this class of contact models.

We note that rigid-bodies constitutes a large class of systems of interest in robotics. However, manipulation of deformable objects is an important area of research interest that is not well suited to the approach in this chapter. Due to the high-dimensional parameter space and state representations of these systems, large scale nonlinear optimization with contact is difficult. This optimization is akin to solving a finite-element mesh for contact forces and parameters which can be computationally expensive and sensitive to small state perturbations.

In this chapter, we focused on contact resolution techniques using LCP. These

approaches are just one of several models available. In the future we plan to evaluate the performance of various models such as [5, 24]. Further, We also implicitly assumed that the geometry of the system and environment is known precisely. Uncertainty in the geometry of the problem however presents interesting and revealing challenges and is another exciting direction to explore. Finally, we formulated a non-linear optimization program (due to the constraints). The characterization of the space of minima (local) may be NP-hard as it will require the solution to a complementarity problem for each set of parameters and forces. The pivoting techniques (such as Lemke) provide efficient solvers but are worst case exponential. A better characterization of the solution space may be possible due to the convex cost. This characterization remains for future work.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Empirical Contact Model Predictive Performance Evaluation

3.1 Goal

In the previous chapter, we derived an identification formulation for rigid-bodies undergoing frictional interactions. The fidelity of the identification depends on the quality of the predictions of the models we used. In this chapter, we present an empirical evaluation of six commonly used contact models in robotics. Specifically, we discuss:

- Contact model predictive performance comparison on the planar impact task,
- Contact model parameter identification variance,
- Model failure modes.

We compare models uniformly by using the underlying physical constraints, referred to as the Energy Ellipse [17]. We further use this representation as a tool in model diagnosis.

3.2 Introduction

Despite the variety and wide usage of existing rigid body contact models, relatively little empirical evaluation and comparison among models has been conducted. The performance of such models has not been quantified; consequently, the limitations of these models are not well understood. Additionally, systematic parameter estimation methodologies have not been developed for the class of rigid impact models. This chapter contributes the first empirical comparison of complementarity amenable rigid-body contact models in the literature.

To perform the empirical comparison, we first review rigid-body contact constraints. From these constraints, we arrive at the Energy Ellipse as discussed in [38]. We use this representation as the basis of the comparison, and show how each model is a parameterized subset of the constraints outlined by the Energy Ellipse. We then perform identification and show the empirical performance of the models on the task.

In the ensuing discussions and results, we show:

- model parameter identification variances,
- fundamental model predictive limitations,
- a characterization of the types of contact and regions of the pre-contact state-space for which the models perform well and poorly.

3.3 Problem Formulation

Rigid contact models couple the Newton-Euler equations of motion to contact equations (non-interpenetration constraints, Coulomb friction, stiction, *etc.*) to predict forces imparted to the bodies, and, subsequent motion by extension. The class of rigid contact models makes three simplifying assumptions:

1. bodies undergoing contact sustain negligible deformation (coupling this assumption with convex geometry implies that a single point of contact need be considered),

2. all impacts and their resolutions are instantaneous, and therefore velocity changes discontinuously,
3. the configuration of impacting bodies does not change during impact (follows from the instantaneity assumption).

These assumptions lead to point or line contacts. The reason that only point and line contact are considered is nuanced. First, we have no computational means of computing contact forces distributed over surfaces. Second, even if we had such a computational tool, general indeterminacy in the rigid contact model implies that it wouldn't represent the "true" solution anyway. Put another way, even if you sample contact points with increasing resolution, you are not guaranteed to converge to a single result.

Consider Fig. 3-1, where a planar object makes contact with a fixed horizontal surface. The configuration of the object is $\mathbf{q} \equiv [x \ y \ \theta]^T$, the velocity $\mathbf{v} \equiv \dot{\mathbf{q}}$, the point of contact $\mathbf{r} \equiv [r_x \ r_y]^T$, and the mass and second moment of inertia m and I respectively. Unless stated otherwise, all points and vectors defined in a "world" (stationary) reference frame. We can represent a rigid model for contact of a planar moving body with a fixed surface as a mapping:

$$\mathbf{f}_c : \{\mathbf{q}, \mathbf{v}^i, \mathbf{r}, m, I\} \rightarrow \{\mathbf{v}^f\} \quad (3.1)$$

where the superscripts i and f denote the pre- and post-contact values (recall that the rigid body assumption, which implies instantaneity, entails that $\mathbf{q}^i = \mathbf{q}^f$, so we omit such subscripts from \mathbf{q}). Under the rigid-body assumption, Newtonian mechanics relates the pre- and post-contact states by:

$$\mathbf{M} \frac{\mathbf{v}^f - \mathbf{v}^i}{dt} = \mathbf{J}(\mathbf{q}, \mathbf{r})^T \mathbf{f}_c \rightarrow \mathbf{v}^f = \mathbf{v}^i + \mathbf{M}^{-1} \mathbf{J}^T \mathbf{c} \quad (3.2)$$

where $\mathbf{M} \in \mathbb{R}^{3 \times 3}$ is the generalized inertia matrix of the rigid body, $\mathbf{c} = \mathbf{f}_c dt$ denotes the impulse imparted to the bodies due to contact, and $\mathbf{J} \in \mathbb{R}^{3 \times 2}$ denotes the Jacobian matrix that maps generalized velocities to contact-frame velocities and contact-frame




Figure 3-1: Planar object making contact with a horizontal surface. The contact point is denoted by C .

impulses to generalized impulses (via the transpose of the Jacobian). Note that this relation holds for the moment before and after contact, but is insufficient to determine the contact impulse (the vector operations above yield three equations and six unknowns: \mathbf{v}^f and \mathbf{c}).

Given the unknown imparted impulse we can solve for the post-contact velocity; choosing the imparted impulse is precisely the job of the contact models. To gain an understanding of how contact models choose \mathbf{c} , in sec. 3.3.1 we will briefly discuss the fundamental constraints of contact. This understanding will lay the foundation to the system identification formulation as well as the data-driven methods.

3.3.1 Fundamental constraints of contact

Under the rigid-body assumption, all contact must obey three fundamental constraints:

- only compressive impulses may be applied along the surface normal (contacts are not permitted to pull),

- impacting bodies must not penetrate,
- the total system momentum must be conserved.

To obtain a visual representation of the feasible impulse space for 2D contact we adopt the treatment discussed in [17]. We write the change in momentum of the object due to contact as:

$$\mathbf{v}^f - \mathbf{v}^i = \mathbf{M}^{-1} \mathbf{J}^T \mathbf{c} \quad (3.3)$$

We write the change in velocity of the contact point by pre-multiplying both sides of the expression by \mathbf{J} :

$$\mathbf{J}(\mathbf{v}^f - \mathbf{v}^i) = \mathbf{J} \mathbf{M}^{-1} \mathbf{J}^T \mathbf{c} \quad \rightarrow \quad (\mathbf{v}_C^f - \mathbf{v}_C^i) = \mathbf{M}_C^{-1} \mathbf{c} \quad (3.4)$$

Here \mathbf{v}_C is the 2×1 velocity vector of the point of contact c , $\mathbf{M}_C^{-1} \equiv \mathbf{J}^T \mathbf{M}^{-1} \mathbf{J}$ is the 2×2 *contact space compliance matrix* (making \mathbf{M}_C the *contact space inertia matrix*). We note that the kinetic energy of the system post-contact may not exceed that of the energy prior to contact; we may write this constraint as:

$$\mathbf{v}_C^{fT} \mathbf{M}_C \mathbf{v}_C^f = \alpha \mathbf{v}_C^{iT} \mathbf{M}_C \mathbf{v}_C^i, \quad 0 \leq \alpha \leq 1 \quad (3.5)$$

Recall that the position of the object does not change during an instantaneous impact, so potential energy remains the same. We re-write this expression using (3.4), and after expansion and substitution:

$$(\mathbf{c} + \mathbf{M}_C \mathbf{v}_C^i)^T \mathbf{M}_C^{-1} (\mathbf{c} + \mathbf{M}_C \mathbf{v}_C^i) = \alpha \mathbf{v}_C^{iT} \mathbf{M}_C^{-1} \mathbf{v}_C^i \quad (3.6)$$

The representation of the expression above is that of an ellipse in 2D due to the positive definite nature of \mathbf{M}_C ; the term $\mathbf{M}_C \mathbf{v}_C^i$ simply acts to translate the ellipse $\mathbf{c}^T \mathbf{M}_C^{-1} \mathbf{c}$ in \mathbf{c} . The union of all ellipses for all $0 \leq \alpha \leq 1$ gives the Energy Ellipse [17].

Fig. 3-2 depicts an Energy Ellipse, where the shaded region is the admissible region based on the fundamental constraints of rigid contact. Two important features

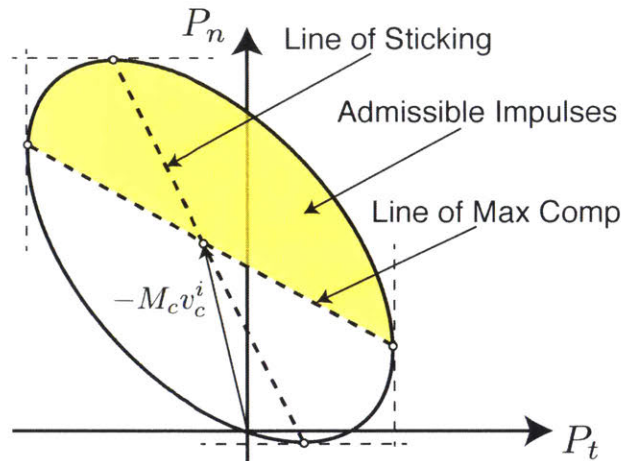


Figure 3-2: The Energy Ellipse: The ellipse is defined in the contact frame where the ordinate and abscissa are the normal and tangential impulses, respectively. The interior of the ellipse demarcates the set of impulses that conserve kinetic energy before and after contact. Contact impulses above the *Line of Maximum Compression* prevent penetration between the bodies making contact. Impulses on the *Line of Sticking* result in zero tangential velocity during contact (sticking contact). The impulses in the shaded region satisfy the constraints of rigid-body contact models.

of the Energy Ellipse are the lines of sticking and maximum compression.

The line of maximum compression delimits the region for which the separation velocity of the contact point changes sign; regions on and above this line obey the fundamental contact constraints. Points along the line of sticking signify contacts for which the relative tangential velocity of the contact point is zeroed due to friction. The intersection of these two lines is the impulse imparted that eliminates all relative velocity between the bodies at the contact point.

The diagram of the Energy Ellipse graphically shows that the three fundamental constraints governing the impact are not sufficient to select an impulse to resolve the impact. Contact models introduce additional assumptions and constraints to choose a unique point within the feasible space. In the next section we briefly introduce the models considered in this paper, and demonstrate how they select points within the Energy Ellipse.

3.3.2 Rigid contact models studied in this paper

The purpose of a rigid contact model is to select a point in the feasible impulse space demarcated with the Energy Ellipse. In this section, we very briefly introduce the six models we used in this study; and refer the interested reader to [28] for further details. The six models studied are all parameterized by two variables, denoted as (μ, ϵ) and often referred to as the coefficient of friction and restitution, respectively. When the contact-space compliance matrix is diagonal, the former is used to regulate the magnitude of tangential momentum of the contact point and the latter is used to regulate the normal component.

The first two models are **Anitescu-Potra Newton** [4] and **Anitescu-Potra Poisson** [4], which are closely related to Stewart-Trinkle [79], differing primarily in the restitution used for contact resolution as their names suggest. The **Drumwright-Shell** [25] also uses a Poisson-type contact model that computes contact forces that maximize kinetic energy dissipation. **Mirtich** [62] is an incremental collision model and a computational implementation of Stronge’s work [81] relating physical work done in the compression and restitution phases of an impact. **Wang-Mason** [87] relates the magnitude of the impulses in the compression and restitution phases of an impact using a Poisson-type restitution. **Whittaker** [89, 44] is an algebraic model that requires solution to a nonlinear system of equations and both 2D and 3D versions have been studied.

We now define two post hoc models, the **Best Post Hoc** and the **Ideal Rigid-Body Bound (IRB Bound)**. The Best Post Hoc model selects the best predicting rigid-body contact model, of the six considered in this paper, after each impact event, and the IRB Bound selects the best predicting impulse from the Energy Ellipse. We emphasize that the IRB Bound is able to select any possible feasible impulse, but the Best Post Hoc only selects a model (spanning a sub-space of the Energy Ellipse). Consequently the IRB Bound quantifies the absolute best performance feasible for any possible rigid-body contact model, and the Best Post Hoc quantifies the best possible performance for the models studied. Neither model is suitable for making




Figure 3-3: Regions of the admissible contact space covered by each model. Note the similarity between the models and the limited coverage of the admissible space.

predictions, since each makes their choice post hoc.

Fig. 3-3 depicts the admissible contact regions for each model, grouped based on graphical similarity. This similarity is not necessarily indicative of the way the regions are constructed. These regions are the union of all predicted impulses for all admissible values of the model parameters.

Fig. 3-3 (a) shows that the Drumwright-Shell and AP Poisson models cover very similar regions of the space. Fig. 3-3 also provides intuition for the definitions used in contact parameters. For example, the coefficient of restitution regulates the magnitude of the normal impulse; therefore, models with vertical side boundaries use the ratio of the post to pre-contact normal momenta. Fig. 3-3 (b) depicts the regions from the AP Newton, Wang-Mason, and the Mirtich models. We note that the definition of the coefficient of restitution used in these three models relates the relative normal velocities pre- and post-contact, resulting in fairly similar regions. Fig. 3-3 (c) depicts the Whittaker model, for which the horizontal boundaries of the feasible region imply that the coefficient of friction relates the tangential momenta pre- and post-contact, differing significantly in interpretation from the previous models.

Two key observations are:

- no single model is able to span the full admissible region of contact (we further note that the regions covered by the various models are quite similar in shape)
- any given model will make predictions only up to the line of sticking, but not beyond it ; therefore the models are not able to predict "back-spin" (i.e., the

tangential velocity of the contact point post-contact changes direction).

The latter issue is sometimes alleviated by contact models that utilize an additional tangential restitution, but the physical interpretation and utility of such restitution is debatable [17].

We note that once an impulse is chosen, we use Newton’s equations of motion to compute the post contact velocity of the center of mass, under the assumption of rigidity.

3.3.3 Model identification and evaluation

System identification applied to inertial parameter estimation of robotic systems has received considerable attention within the robotics research community [29, 30, 35, 47, 52] and several practical implementations have been developed and applied to industrial robots [65]. Control architectures have also been developed to estimate and adapt to uncertainties over these dynamic parameters [42, 76], and have been shown to exhibit asymptotic convergence to desired trajectories in the case of free space motion of manipulators.

System identification is a necessary step to maximize the predictive performance of these models, yet to our knowledge no systematic formulation of this step has been proposed. To this end, we first propose a system identification formulation in impulse space, leveraging the Energy Ellipse. Next, we briefly discuss the experimental setup and data collection protocols used in this study; we then apply the system identification procedure to find the optimal parameters for each of the studied models, allowing us to empirically compare predictions and use the IRB Bound and Best Post Hoc models to establish upper limits in predictive performance. Finally, we study individual experimental trials and demonstrate the large variability in model parameters across experimental trials (indicating challenges in interpreting parameters). We also quantify the predictive range of each model by evaluating the number of outcomes across individual instances of the experiment that can be explained by *any* choice of parameter set.

To perform system identification, we first compute the impulse imparted on the object from pre- and post-contact states ([29]) under the ideal rigid body impact model by solving:

$$\mathbf{c}_{irb} = \arg \min_{\mathbf{c}} \|\mathbf{M}(\mathbf{v}^f - \mathbf{v}^i) - \mathbf{J}\mathbf{c}\|_2 \quad (3.7)$$

Then we can identify the parameters for a single impact by solving the following optimization program:

$$(\mu^*, \epsilon^*) = \arg \min_{(\mu, \epsilon)} \|\mathbf{c}_{irb} - \text{Model}(\mu, \epsilon)\|_2 \quad (3.8)$$

$$\text{s.t.} \quad 0 \leq \mu \leq \mu_s \quad 0 \leq \epsilon \leq 1 \quad (3.9)$$

The optimization program is convex within the constraints (as will be explained in more detail in Sec. 3.4) and takes place in the space of imparted impulses and the Energy Ellipse. The upper bound on the coefficient of friction μ_s is the lowest value of the coefficient of friction for which the model predicts sticking contact and can be found with a line search using a bisection method. For values of $\mu > \mu_s$ and a fixed ϵ , the predicted impulse from the model will remain the same and the cost function will not have a gradient, making the optimization ill-posed numerically. This implies that pre and post contact conditions exist which do not provide sufficient information to estimate μ correctly and render this parameter unobservable, more on this in Sec. 3.4. To perform system identification for a batch of impacts, we simply sum up the cost function over the set of impacts (Sec. 3.4).

3.3.4 Experimental data

In order to evaluate and compare the contact models we conceived the experimental setup depicted in Fig. 3-4 [28, 29]. The purpose of the apparatus is to autonomously collect measurements of the trajectory of falling planar objects initialized from a specifiable set of initial conditions. We chose this experiment because it captures key features of a dynamic frictional interaction while being amenable to careful monitor-




Figure 3-4: Experimental setup: Autonomous data collection for planar impact data. A planar object is constrained to move between two lubricated vertical glass planes. The robot uses a directional magnetic latch to pick up the object across the glass and provide control over position and velocity of the object. The Vicon motion capture system is used to track the states of the object.

ing and instrumentation resulting in high quality measurements; for further details see [28].

The experimental setup is composed of a “dropping arena” designed to constrain the motion of planar objects to a plane using lubricated glass planes. The planar objects are equipped with markers to be tracked using the Vicon motion capture system at 250 Hz, and directional magnets allowing the robot to latch onto the object and prescribe desired initial state to the object for each drop.

We collected a total of 2,000 drops. In some instances the tracking system may lose frames or suffer spurious errors due to reflections; to detect and remove these effects each experimental run was subjected to a series of tests to validate the correctness of the measurements. The three most important tests were: i) the energy test: verifies constant energy during ballistic motion and energy loss during times of contact, ii) a frame drop test, and iii) a deviation test: checks deviations of trajectories during ballistic motion from parabolic paths. A total of 1,718 of the 2,000 drops passed all tests, and we used this set for the rest of this study. An example of the recorded configurations for one experiment is shown in Fig. 3-5. Fig. 3-6 shows the potential, kinetic, and net energies of the object for a snippet of the trajectory as used for test (i) above.




Figure 3-5: The configurations of the object as a function of time during an experimental run. θ is in radians but has been scaled by the radius of gyration of the object.




Figure 3-6: Example test of object energy during the first two bounces. The net energy of the object is approximately constant between two impact events. The net energy drops after each impact event.

To detect the time of contact we use the fact that the object undergoes ballistic motion, consequently the acceleration of the object is smooth. Any impact event induces a spikes in the second time derivative of the state vector, and we use a threshold on these spikes to detect events. We further use the ballistic motion of the object to fit parabola to the pre- and post-contact trajectories, thereby mitigating the effects of noise on the state measurements and yielding better approximations to velocities of the object; for more details see [28].

3.4 Results

In this section we present the results of the application of the system identification procedure (described previously in Sec. 3.3.3) using the experimental data (collected as described in Section 3.3.4). We first select N random trials to find the parameter set (the Ensemble (μ, ϵ) parameters) that best describe these trials for each model, and we evaluate the performance of the models in predicting the outcome of impacts on previously unseen data. We next compute the model parameters, (μ_i, ϵ_i) that best describe each individual drop (the individual best parameters); this computation captures the variance in the parameters and will highlight some of the limitations of the contact models.

Ensemble parameter identification

Ensemble parameters were identified from a randomly selected subset of the data as per the procedure described in Section 3.3.3. The topography of the cost function is depicted in Fig. 3-7 (a). The topography of the resulting function is convex, which is to be expected since the region of admissible impulses predicted by all models is convex and the values of the predicted impulses grow proportionally to the magnitudes of the parameters. Intuitively, if there exists a pair of parameters for a given that perfectly predicts the impulse, then the cost associated with that pair would be zero and all other points within the convex set would have error values that would grow with distance away from the optimal pair. For a point outside the set of admissible

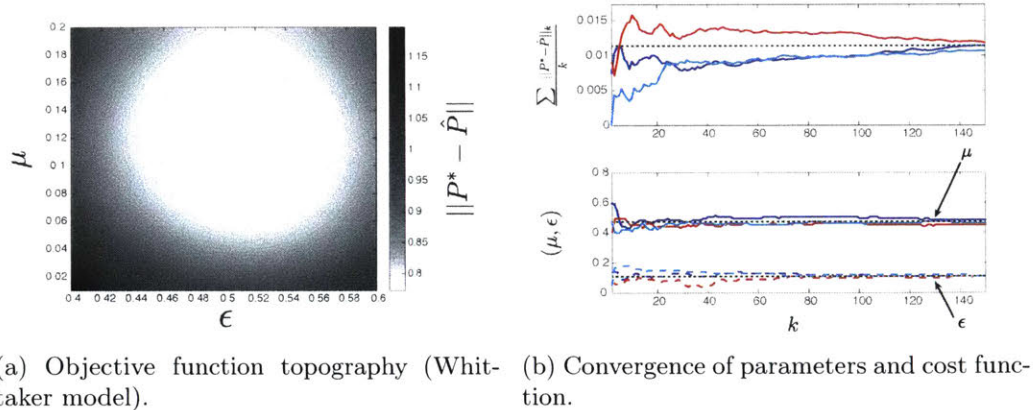


Figure 3-7: The left panel shows the topography of the cost function plotted as a function of the parameters for the Whittaker model where the summation is over 150 drops and lighter colors denote better objective values. The right panel shows the convergence behaviour of the model parameters and objective function values for 3 randomly selected subsets of data, plotted as a function of number of data in the subset k .

impulses, due to the imposed constraint $\mu \leq \mu_s$, then there is a point that lies on the boundary of the set for which the distance takes on a minimum; this observation arises due to the convexity of the set.

Fig. 3-7 (b) shows the convergence of the parameters as a function of the number of data used, we selected a random set of trials and observed that the parameter values converge to their final values within approximately 100 trials. We repeated the process several times with different randomly selected trials to validate the convergence behavior, three of which are shown in the figure. This process was repeated for all models and Tab. 3.1 presents the results.

As can be seen from Tab. 3.1, the values for the coefficient of friction found for each model are much lower than the values estimated from an inclined surface experiment between PLA (material of the object used) and Aluminium (material of the impact plate) (0.3). We believe that the major contributing factor to the discrepancy is the much more complicated interaction that the coefficient of friction tries to model in the highly dynamic dropping experiments. The object potentially undergoes a number of rich and difficult to model interactions in a very short period of time such as small local deformations, finite-time impact, vibrations, and impulsive torsional frictional

Table 3.1: Average-best parameter identified values and the percentage a model was chosen as the best and worst model

Model	μ	ϵ	Best%	Worst %
DrumShell	0.081	0.516	4	11
AP Poisson	0.101	0.526	7	6
AP Newton	0.084	0.547	25	0
Mirtich	0.062	0.558	17	43
Wang-Mason	0.120	0.537	16	10
Whittaker	0.111	0.484	31	30

forces in conjunction with linear ones. The use of just one parameter, which in this case is state independent, may not be sufficient to describe these rich interactions. In contrast, the inclined surface experiment is slow and controlled, and one parameter is sufficiently descriptive to explain the observations. Further, [17] discusses the simplicity of these models, in that they attempt to capture a complicated phenomena in a minimalistic way and with intuitive parameters. This intuition is valuable in that it gives a direct interpretation to the way these parameters affect the outcome of a prediction, but it is unreasonable to expect the values found for these parameters to match those found from other experiments. The coefficient of friction in this context is not necessarily the same as that measured from the inclined surface experiment, though they have the same name. It is important to note that using the inclined surface experimental value for μ will lead to significantly poorer results.

Fig. 3-8 depicts the estimated PDF of the ℓ_2 -norm error of predicted v_o (we normalize the angular velocity by the radius of gyration so that the linear and angular terms are of the same order). The “Best Post Hoc” model is also plotted to show the performance upper bound on the analytical model; that is to say, if, for each trial, we knew *a priori* which model would perform best and used that model to make the prediction, this is the best performance we could expect (using this set of analytical models). Tab. 3.1 shows the number of times each model was chosen as the best and worst performing for the data-set. We can see that the Whittaker model is quite erratic, while the most consistent model is the AP Newton model, with a significant number of best predictions and almost no worst.




Figure 3-8: Estimated probability density functions (PDFs) of the ℓ_2 -norm error in post-contact velocity of the center of mass.

The “Ideal” contact model is an oracle model that computes the optimal contact impulse to predict the post-contact velocity vector under the ℓ_2 -norm. It does this by choosing an impulse from the permissible region in the energy ellipse and has access to the pre- and post-contact states. The key observation here is that even this oracle model has a non-zero error in predicting outcomes of these contact events.

Recall that the contact models select $P_{2 \times 1}$ and use the contact Jacobian $J_{3 \times 2}$ to compute $v_{o,3 \times 1}$ but if we allow for a torsional component, i.e., $P_{3 \times 1}$, we may be able to improve predictions at the cost of relaxations of assumptions of rigid-body contact models and constraints.

It is important to note that the error introduced by the rigid body assumptions is not distinguishable from error resulting from uncertainty in sensor measurements, but our setup was designed to minimize the latter errors. The implication of this result is that it may be impossible to construct a zero error model without violating the rigid body or rigid contact models, meaning that the best we can hope to achieve in terms of performance is the “Ideal” model.

We can further examine the errors made in the models’ predictions to quantify regions of the input space in which we expect models to perform comparatively better and worse. This information is potentially very valuable if we are interested in select-


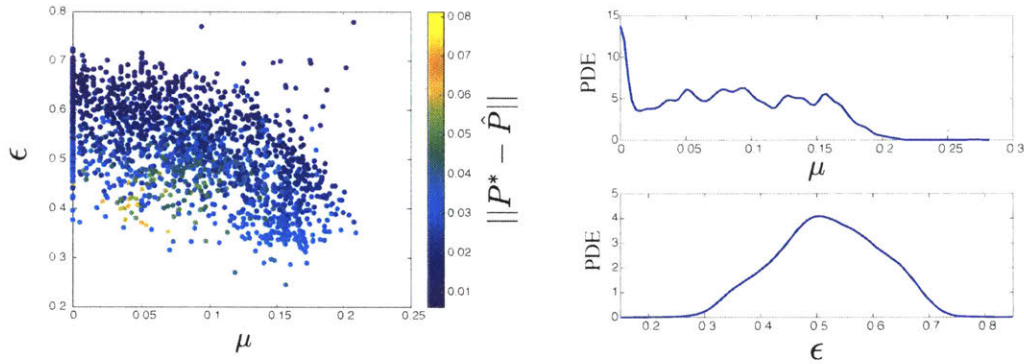


Figure 3-9: The performance of the model depicted as a function of the normal and tangential components of the pre-contact momentum of the system measured at the contact point. Deep blue regions represent high accuracy, whereas yellow regions predict high error regions.

ing initial conditions that are easier to predict or rejecting initial conditions that may be more difficult to predict (for purposes of control). Fig. 3-9 depicts the performance of the AP Poisson model in predicting the outcome of an contact (error in predicting the velocity of the COM) as a function of the normal and tangential components of the pre-contact system momentum measured at the contact point. The region in which the tangential momentum at the contact point is near zero or changes direction proves to be the most difficult to predict; we will discuss this phenomenon in Sec. 3.4. We can demarcate reliable and predictable regions of the feature space using k-Nearest neighbors or other clustering algorithms.

Individual best parameter identification

The optimization we performed in Section 3.4 to find the ensemble parameters hides some of the details and idiosyncrasies of individual trials. In this section we perform the identification procedure for each individual drop, i.e., we find the values of (μ, ϵ) that best explain each drop. Fig. 3-10 a) shows the values of these parameters colored by performance for the AP Newton model and b) shows the distribution of the



(a) AP Newton model parameters (μ, ϵ) identified for each drop, color-coded by error where from blue to yellow performance worsens. (b) Probability density estimates of (μ, ϵ) over all data for the AP Newton model. Note the wide distribution in both parameters.

Figure 3-10: Identified values of (μ, ϵ) for each drop, where a) shows the values color-coded by the ability of the identified pair to predict the impulse, and b) showing the distribution of the parameters of the data-set.

identified parameters.

Studying Fig. 3-10 reveals the following five important insights:

1. A sharp peak (concentration) of $\mu = 0$ is observed in both figures, which we will explain shortly.
2. The lower left region in Fig. 3-10 a) clearly shows a loss in performance, we will also explain shortly.
3. The distribution of ϵ is uni-modal with a peak centered around 1.
4. When compared to the ensemble identification procedure we see that the means of the distributions of μ and ϵ are very close to the optimal values computed for the ensemble parameters.
5. Neither distribution is tightly localized about a single value. This is perhaps the most important feature to be noted since it is common practice to select a single value of the parameters. A single value for ϵ may be chosen by fitting a Gaussian to the distribution and taking its mean but deciding on a value of μ is much more difficult, particularly given that the values are significantly smaller

than the values measured during a standard inclined surface experiment and that the distribution is relatively flat.

The fact that such variation in parameter values exists, particularly for the value of μ , begins to highlight issues of interpretation of the parameters. The intuitive expectation for a coefficient of friction is that it should be constant for all the impacts, and the values should be localized to very similar values, but these experiments demonstrate that this is not the case. It is our belief that the coefficient of friction in these models should be treated as more of a restitution variable that regulates the magnitude of the tangential impulse as a function of the pre-contact conditions rather than the notion of coefficient of friction we ascribe to the inclined surface experiments.

The anomalous cases for which $\mu = 0$ can be isolated and studied; Fig. 3-11 a) shows an example. When the contact point lies directly below the center of mass of the object, the Energy ellipse's major axis aligns with the normal impulse direction and the convex set of predictions for the outcome of the contact collapses to a very narrow band, resulting in a very small predictive range. A small prediction range coupled with sensor noise leading to errors in localization of the contact point severely limits the prediction range of the models. In the case shown in Fig. 3-11 the space spans to the left but the measured impulse lies to the right; consequently, the value of the coefficient of friction tends to zero to get as close as possible to the target value. This issue highlights the difficulties in interpreting the coefficient of friction within the context of these contact models.

The second set of anomalous cases is characterized by the inability to predict "spin-reversal", an example of which can be seen in Fig. 3-11 b). In these scenarios the measured impulse does not fall within the predictive range of the model, and a residual error will always be present, explaining the region in Fig. 3-10 where the parameters take on values with non-zero errors. Tab. 3.2 quantifies the fraction of data that lie within the convex set of each model, and the relative ratios are quite low. This indicates that the models suffer from potentially restrictive prediction spaces. Note that models with similar regions in the energy ellipse exhibit similar ratios, exemplified by the AP Poisson and the Drumwright-Shell models.




Figure 3-11: a) Example of $\mu = 0$, this impact is characterized by the point of contact being located directly under the center of mass of the object. In this situation the Energy ellipse is axis aligned to the normal impulse direction. b) Example of a measured impulse outside the convex set of predictions made by the model due to “spin-reversal”.

Table 3.2: Fraction of drops that lie within the predictive range of each model

Model	Fraction
DrumShell	0.355
AP Poisson	0.357
AP Newton	0.458
Mirtich	0.413
Wang-Mason	0.455
Whittaker	0.456

These issues are, in part, responsible for the “bumps” in the pdfs of the error plots in Fig. 3-8. In particular, the challenge in predicting back-spin and impulses that lie outside the predictive range of the models leads to the tail-end errors of the distribution.

3.5 Related Work

Rigid-body contact models are widely used in simulators. Many of these simulators serve as tools for robotics researchers. Many contemporary algorithms rely on accurate models of the dynamics of a robot as well as models of interactions of the robot with its environment ([19, 42, 71]). Dynamic frictional interaction between

bodies is important to robotic systems undergoing contact, yet has proven to be a challenging phenomenon to model. Challenges include but are not limited to non-smooth dynamics (i.e., transitions from none contact to sticking/sliding contact), near-impulsive forces with large magnitudes, local deformations (difficult to predict elastic/plastic indentations), wave propagation through the bodies (e.g., vibrations along rods from impact), sensitivity to initial conditions, and effecting friction—a complex phenomenon that emerges from the interactions between microscopic surface asperities—through necessarily coarse models. Despite these difficulties various contact models have been proposed and generally follow one of three paradigms: i) fully deformable, elasto-dynamic models; ii) pseudo-rigid, compliant models; and iii) rigid contact.

The elasto-dynamic models allow deformations for contacting bodies over a finite contact time, can approximate elastic/plastic deformations, and are able to model vibrations along the object. These modeling capabilities come at a high computational cost, large number of variables to tune to each contact pair, and much of their predictive power is lost due to high sensitivity to uncertainty of initial conditions. As such these models have not seen much prevalence in robotics, where speed and simplicity play an important role and necessary system identification cannot yet be conducted accurately in real time.

The class of pseudo-rigid, compliant contact models (often designated “penalty methods”) is the oldest, beginning with the seminal work of Heinrich Hertz [41] followed by numerous authors such as [53, 55, 59, 77]. These models use a rigid core and a thin, deformable layer and apply contact forces proportional to the magnitude of deformation. Simulation software such as [64] (open-source) and [56] (proprietary) have been developed under this paradigm, and akin to this modeling paradigm software such as [84] have been developed that use relaxation to soften contact with rigid-body contact models.

The validity of the predictions of such models is still a matter of debate, allowing penetration and permitting contact forces proportional to the magnitude of the penetration smooths out the dynamics of the bodies undergoing contact yet may not be

consistent with physics, especially for rigid-body dynamics. Furthermore, [50] shows that the predictions of these models do not match experimental observations from the studied prehensile manipulation maneuver, these models are not able to distinguish between sticking and sliding interactions and predict penetration when there should not be any.

Rigid contact models resolve contact without permitting geometric intersections between bodies. These methods assume that deformations are nonexistent and contacts “appear” discontinuously as a function of configuration—meaning that impacts must be resolved instantaneously. The rigid contact model distinguishes between non-impacting and impacting contacts; while the combination of complementarity conditions, Coulomb friction, and stiction constraints often yields a fully determinate model for non-impacting contact, impacting contact is underdetermined. Rigid impact models effectively add constraints (like restitution) to select points in the feasible impact space (a concept described by [17]). Rigid-body impact models span centuries of research, starting with Newton, then seeing contributions from Poisson, and then [89, 81, 87]. These models typically require 2-3 parameters (often referred to as normal restitution and friction coefficient) and are commonly used in robotic manipulation [19, 42] and locomotion [71].

Despite the variety and wide usage of existing rigid impact models, little if any empirical evaluation and comparison among models has been conducted. The performance of such models has not been quantified, consequently, the limitations of these models not well understood. Additionally, systematic parameter estimation methodologies have not been developed for the class of rigid impact models. This chapter is one of the first attempts at filling this gap in the literature.

The focus of this chapter was on the complementarity formulations of rigid-body contact models with two principle parameters (friction and restitution). We did not study the pseudo-compliant models or other friction models such as the Stribeck, the Dahl, or the LuGre models [60]. A comparison against these models would be an interesting line of research.

3.6 Discussion

Rigid contact models are popular due to their simplicity and computational efficiency yet they exhibit clear limitations in predictive power and accuracy. In this chapter, we empirically demonstrated:

- predictive range limitations,
- challenges with system identification,
- difficulty of interpretation of the model parameters,
- contact model failure modes (singular impacts and inability to predict back-spin).

The limitations of these analytical models can partially be attributed to the parameterization of a subset of the admissible impulse space. The independence of the parameters from the inputs to the models and the limited regions for prediction (Tab. 3.2) contribute to the errors in predictions. The rigid models studied assume a single point contact and instantaneous impact. This assumption is only an approximation and the impact event occurs over a finite time and involves a small amount of local deformation leading to a distributed impact impulse over a patch. Using a point contact means these models can only apply a linear impulse, when an impulse wrench may be a better approximation to the interaction.

Two approaches can remedy these shortcomings, we may either learn the mapping directly from the impulses and do away with the parameterization all together (data-driven approach) or use the model predictions as a prior and correct the impulse predictions using a data-driven model (data-augmented approach). In the next chapter, we explore these approaches and compare performances across analytical and learned models.

An open challenge we did not address in this chapter is uncertainty propagation through the contact models, i.e. given a distribution over the input space (states and/or parameters), what does this distribution over post-contact states look like?

We are unaware of techniques other than particle filters for this purpose. The hard-nonlinearity (discontinuity) in the post contact states biases unscented Kalman filters and introduces challenges in smooth filtering techniques. Given computationally tractable uncertainty propagation, we may alleviate some of the challenges in these contact models, though we would still not address the predictive range limitations.

This study demonstrates the importance of data in model validation and highlights the importance of experimentation in parameter selection, and its potential role in developing data-driven contact models [33, 91]. We hope that this study serves as motivation to experimentally validate other models frequently leveraged in robotics and leads to their better understanding.

Chapter 4

Sample-efficient Residual Learner

Contact Models

4.1 Goal

In this chapter, we address some of the shortcomings in point contact model predictive performances discussed in the previous chapter. We demonstrate a residual learner formulation that corrects the predictions of these models in the contact frame.

This formulation provides a principled approach to learning corrective models that are:

- sample-efficient,
- generalizable,
- provide a mechanism for relaxing analytical model constraints,
- uncertainty quantification.

These benefits are afforded by virtue of corrections at the impulse level, where we correct the low-dimensional contact impulse – the cause, as opposed to the post-contact state – the effect.

4.2 Introduction

The contact models we studied in the previous chapter are coarse approximations to complex frictional interactions. The assumptions employed in their construction afford significant computational speed-ups at the cost of accuracy.

We have learned one important limitations of these models due to their construction: limited predictive range (within the Energy Ellipse). Further, we know that no contact event is perfectly rigid or instantaneous. The central question in this chapter is then how do we address these restrictions without resorting to higher-order models that incur significant computational burden.

These models compute a linear impulse in the contact frame that satisfies rigid-body constraints. Given the correct impulse, and assuming the contact event is perfectly rigid, then the post-contact state should be predicted with high-accuracy.

However, when fitting parameters, we optimize the analytical models for best average performance. Further, some outcomes lie outside the model predictive ranges (for any choice of parameters). Our key insight is to correct the predicted impulse of the contact model before computing the post-contact states. We can relax the rigid-contact assumption by learning a residual that accounts for an impulsive wrench rather than just a linear term (see following sections for details).

Intuitively, since the models already provide a good first guess, then learning the residual error is an easier task. Further, by making corrections at the impulse level, we are able to leverage many of the invariances in the mapping from the pre-contact states to the post-contact states. These two points afford significant sample-efficiency and provide a well-posed learning problem which adhere closely to physical constraints.

Critically, the residual model also provides a mechanism to propagate uncertainty. This is a central feature, missing in current contact models.

4.3 Problem Formulation

In this section, we first briefly review the key concepts of contact models from previous chapters. In particular, we revisit the Energy Ellipse and use it as a starting point for the residual learner model. We then introduce two classes of data-driven contact models. We further show how to propagate uncertainty using the data-driven models.

4.3.1 Review of the Energy Ellipse

In chapter three, we discussed the rigid-body contact map for a single-body. From eq. 3.1, we wrote:

$$\mathbf{f}_c : \{\mathbf{q}, \mathbf{v}^i, \mathbf{r}, m, I\} \rightarrow \{\mathbf{v}^f\} \quad (4.1)$$

In words, the input to this mapping is the pre-contact states together with the contact formation and inertial properties. The output is then the post-contact velocity of the object.

This mapping does not have a functional form yet. Rigid-body contact models assume that deformations are negligible during impact, and resolve contact instantaneously. This assumption means that the configuration of the object pre- and post-contact does not change and the impact happens at a single point on the body. Point contacts permit only linear impulses to be transmitted between the bodies. We can write the dynamic equations of motion for the instant before and after the impact for a rigid-body as:

$$\mathbf{v}^f = \mathbf{v}^i + \mathbf{M}^{-1} \mathbf{J}^T \mathbf{P} \quad (4.2)$$

where $\mathbf{M}_{3 \times 3}$ denotes the inertia matrix, $\mathbf{J}_{2 \times 3}$ is the contact Jacobian, and $\mathbf{c}_{2 \times 1}$ is the linear imparted impulse. The rigid-body assumption gives Eq. 4.1 a functional form in Eq. 4.2 but is still unresolvable since \mathbf{c} and \mathbf{v}^f are both unknowns.

The impulse \mathbf{c} is constrained by the law of conservation of energy and that no penetration can occur. To derive a mathematical representation for these two constraints,

we pre-multiply Eq. 4.2 by \mathbf{J} :

$$\mathbf{v}_C^f = \mathbf{v}_C^i + \mathbf{J}\mathbf{M}^{-1}\mathbf{J}^T\mathbf{c} \rightarrow \mathbf{v}_C^f = \mathbf{v}_C^i + \mathbf{M}_C^{-1}\mathbf{c} \quad (4.3)$$

where \mathbf{v}_C denotes the velocity of the contact point, and \mathbf{M}_C represents the effective inertia at the point of contact. We may now express the conservation of kinetic energy for the contact point as:

$$(\mathbf{v}_c^f)^T \mathbf{M}_C \mathbf{v}_c = \alpha (\mathbf{v}_c^i)^T \mathbf{M}_C \mathbf{v}_c^i, \quad \alpha \in [0, 1] \quad (4.4)$$

Plugging in Eq. 4.3 we arrive at:

$$(\mathbf{c} + \mathbf{M}_C \mathbf{v}_C^i)^T \mathbf{M}_C^{-1} (\mathbf{c} + \mathbf{M}_C \mathbf{v}_C^i) = \alpha (\mathbf{v}_c^i)^T \mathbf{M}_C \mathbf{v}_c^i \quad (4.5)$$

Eq. 4.5 defines an ellipse in the impulse space (\mathbf{c}), with an offset of $-\mathbf{M}_C \mathbf{v}_C^i$ which is the incident momentum of the contact point. This ellipse is called the Energy Ellipse and is explained in more detail in ([18, 31]). Fig. 4-1b shows an example of an Energy Ellipse, the line of maximum compression denotes impulses that yield zero separating velocities, and the line of sticking denotes the impulses that lead to a zero relative tangential velocity component. The lightly shaded region is the space of admissible impulses subject to the energy and non-penetration constraints. These constraints are still not sufficient to resolve the contact, contact models introduce additional constraints and parameters to select a unique imparted impulse from this admissible set.

In chapter 3, we showed one approach to estimating contact model parameters through minimizing predictive error:

$$\underset{(\mu, \epsilon)}{\text{minimize}} \quad \sum_i \|\mathbf{v}_O^f - \mathbf{v}_O^i - \mathbf{M}^{-1}\mathbf{J}^T \mathbf{f}_c(\mu, \epsilon)\|_2 \quad (4.6)$$

Intuitively, we compute the optimal pair of contact model parameters (μ, ϵ) to minimize the average error in post-contact velocity. In the next subsection, we build




Figure 4-1: a) planar object making contact with a fixed horizontal surface, b) Energy Ellipse.

on the Energy Ellipse interpretation of the contact models for the data-driven contact models.

4.3.2 Learning rigid-body contact models

In this section, we propose two data-driven contact model formulations. The first class is purely data-driven and uses Gaussian Processes (GPs) trained on a feature space inspired by the analytical rigid contact models. The second class uses data to reinforce analytical models, either by informing parameter choice or correcting the predictions made by the models. The former approach is an instance of weak prior learning, while the latter constitutes strong prior learning given the predictions of the analytical models.

Class 1 - Data-driven contact models

Let's assume we would like to learn the contact map for rigid-body contact models using supervised learning techniques. Following Eq. 3.1, we may write:

$$\bar{f}_c : \{\mathbf{q}, \mathbf{v}^i, \mathbf{r}, m, I\} \rightarrow \{\bar{\mathbf{c}}\} \quad (4.7)$$

$$\mathbf{v}^f = \mathbf{v}^i + \mathbf{M}^{-1} \mathbf{J}^T \bar{\mathbf{c}} \quad (4.8)$$

where we optimize over the accuracy in predicting the post-contact velocity of the center of mass of the object. Why may this not be a good idea for rigid-body contact models?

Using this feature space gives the learned models a large amount of flexibility at the cost of sample complexity due to the 10-D input space (for a single rigid body). We’re not exploiting the structure and assumptions of rigid-body contact to reduce the dimension of the feature space. The key insight here is to leverage the invariances and constraints afforded by the rigid-body contact assumption using the Energy Ellipse.

The Energy Ellipse is constructed from the input space of the contact model using a known transformation. The resulting vector field is spanned by the eigenvectors of the effective inertia at the point of contact \mathbf{M}_C (here 2×2 symmetric and positive definite). The velocity of the contact point \mathbf{v}_C narrows the choice of contact impulse to a specific region. Rather than using the full input vector, we apply the known transformation and use one of the following two feature-set:

$$\mathcal{X}_1 = \{\mathbf{M}_{c,11}, \mathbf{M}_{c,12}, \mathbf{M}_{c,22}, \mathbf{v}_{c,t}, \mathbf{v}_{c,n}\} \quad (4.9)$$

$$\mathcal{X}_2 = \{(\mathbf{M}_c \mathbf{v}_c)_t, (\mathbf{M}_c \mathbf{v}_c)_n\} \quad (4.10)$$

$$\mathcal{Y}_1 = \{c_t, c_n\} \quad (4.11)$$

Here \mathcal{X}_1 and \mathcal{X}_2 are respectively 5 and 2 dimensional input spaces to the learner model. We note the significant reduction from the 10-D space of the original input vector. \mathcal{X}_1 attempts to capture the key features required to reconstruct the space of admissible impulses, while \mathcal{X}_2 is a coarse approximation of this space, where we consider the aggregate effect of inertia and velocity as the linear momentum of the contact point. We refer to these two models as “**Data-driven Rigid Contact Models**” due to their similarity to the analytical models.

The data-driven rigid contact models inherit the same properties as the analytical models, they assume a point contact that occurs instantaneously which results in an instantaneous linear impulse. Real objects often undergo small deformations during




Figure 4-2: The new Energy Ellipsoid due to the addition of the contact wrench to account for deformations and finite contact time.

contact over a finite time horizon were forces are distributed over a finite area.

We can account for a small amount of deformation and finite contact time by relaxing the requirement of using only a linear impulse. By allowing an instantaneous wrench ($\mathcal{X}_1, \mathcal{Y}_2 = \{c_t, c_n, \tau\}$), we afford the model greater expressibility with a minimal cost to computation and sample-complexity. We refer to this model as “**Data-driven Contact Model**” since this model deviates from the rigidity assumption, Fig. 4-2.

To learn these contact models we use Gaussian Processes (GPs) ([72]) with an Automatic Relevance Determination (ARD) Square Exponential Kernel. We use this kernel to compensate for the relative differences in magnitudes of the features. This is an important detail, since the normal impulse component tends to be larger than the tangential in the absence of significant angular velocity.

Class 2 - Data-augmented contact models

The models presented in sec. 4.3.2 are able to make predictions that span a larger subspace of the Energy Ellipse than their analytical counterparts, but can conceivably produce infeasible predictions until trained on a sufficient number of data.

In contrast analytical models, by construction, produce feasible predictions for reasonable choices of parameters. We can use data to reinforce these models and to this end we propose two formulations: i) using an analytical model and learning the residual errors in prediction using a GP, and ii) learning the optimal choice of model parameters for an analytical model given pre-impact states. Analogous to the naming convention in sec. 4.3.2 we will refer to these models as “**Data-reinforced Rigid Contact Models**”. The former type of model is amenable to using an instantaneous wrench (i.e. \mathcal{Y}_2) and we will refer to the resulting model as “**Data-reinforced Contact Model**”. In both cases we use the feature space \mathcal{X}_5 and the same kernels as sec. 4.3.2.

4.3.3 Uncertainty propagation

In this section, we show how the residual learned models may be useful to propagate uncertainty. First, recall that the residual learner corrects the impulses predicted by the contact models before passing them onto the dynamic equations of motion, i.e.:

$$\begin{aligned}\mathbf{c}_{cor} &= \mathbf{c}_{\mathcal{M}} + \mathbf{c}_{res} \\ \mathbf{v}^{post} &= \mathbf{v}^{pre} + \mathbf{M}^{-1}(\mathbf{f}_{ext}dt + \mathbf{J}^T \mathbf{c}_{cor})\end{aligned}$$

where \mathbf{c}_{cor} is the corrected impulse and \mathbf{c}_{res} is the residual learner correction impulse. We learned the residual as a supervised learning problem using Gaussian Process regression, i.e.:

$$\mathbf{c}_{res} \sim \text{GP}(0, \mathbf{K}_{xx})$$

For a particular contact event, we draw a sample from the distribution of the

residual learner:

$$\mathbf{c}_{res} \sim \mathcal{N}(\boldsymbol{\mu}_{res}, \boldsymbol{\Sigma}_{res}) \quad (4.12)$$

The corrected impulse distribution is then:

$$\mathbf{c}_{res} + \mathbf{p}_{\mathcal{M}} \sim \mathcal{N}(\boldsymbol{\mu}_{res} + \mathbf{c}_{\mathcal{M}}, \boldsymbol{\Sigma}_{res}) = \mathcal{N}(\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_{res}) \quad (4.13)$$

This distribution is propagated through Newton's second law resulting in a Gaussian distribution for post-contact velocity:

$$\mathbf{v}^{post} \sim \mathcal{N}(\boldsymbol{\mu}_{post}, \mathbf{J}\mathbf{M}^{-T}\boldsymbol{\Sigma}_{res}\mathbf{M}^{-1}\mathbf{J}^T) \quad (4.14)$$

$$\boldsymbol{\mu}_{post} = \mathbf{v}^{pre} + \mathbf{M}^{-1}(\mathbf{f}_{ext}dt + \mathbf{J}^T\boldsymbol{\mu}_w) \quad (4.15)$$

Note that the contact model introduces uncertainty in the post contact velocity but does not change the configuration of the object or its distribution. This is direct result of one of the basic assumptions of rigid-body contact models, that contact is resolved instantaneously and configurations do not change. To capture the expected complex multi-modal distributions, we represent the belief as a Gaussian mixture:

$$p(\mathbf{s}_t) = \sum_{i=1}^m \mathcal{N}(\mathbf{s}_t | \boldsymbol{\mu}_{t,i}, \boldsymbol{\Sigma}_{t,i}) \quad (4.16)$$

where m is the number of mixtures, and $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$ are their mean and covariance, respectively. We then propagate the belief by evaluating the integral:

$$p(\mathbf{s}_t) = \int p(\mathbf{s}_t | \mathbf{s}_{t-1})p(\mathbf{s}_{t-1})d\mathbf{s}_{t-1} \quad (4.17)$$

The selection of a GP model for the residual learner produces a Gaussian mixture posterior distribution for the contact models which complements the Gaussian mixture representation of the belief. We exploit this Gaussian model and mixture to efficiently propagate belief using a hybrid particle Bayes filter. Starting from the

Gaussian mixture, at the start of every dynamic step we take m samples, then propagate the samples until contact, then invoke the data-augmented contact model, and finally update the state belief distribution using the weighted Gaussian mixture, then repeat for the next dynamic step.

This approach is a simplified and slightly modified version of the belief propagation step of the GP-SUM algorithm [10]. The modification we make is to guarantee feasibility of the predictions of the data-augmented contact model. The residual learner does not guarantee that the distribution satisfies physical constraints, like non-penetration and energy balance.

When sampling from the posterior of the contact model it is important to maintain feasibility of the trajectory. In our implementation, a trajectory is infeasible if it results in penetration or an increase in total energy of the system.

The premise of our proposed simulation scheme is that contact free rigid-body dynamics are deterministic and that the main source of uncertainty is contact. As such, contact events are pivotal in propagating and generating beliefs. A “dynamic stochastic step” in this simulation framework is then a transition from post-contact states and beliefs of a system to the next post-contact states and beliefs. This macro-step is a measure between two contact events, and is different from the deterministic simulator time-step. In a deterministic simulator, the time-step drives the simulation forward, and is also used here during the deterministic phases of motion.

4.3.4 Illustrative Example:

For the empirical evaluation of our simulation paradigm we predict the outcome of planar dice rolls. We provide the initial noisy position and velocity of the dice and simulate forward. Fig. 4-3 shows the first two contacts of an empirical dice roll, and overlaid are samples of the states of the dice from the posterior distribution predicted by the stochastic simulation. The **red** samples from the initial state distribution are propagated forward to the first contact event. Post contact, the uncertainty is increased as can be seen by the dispersion in the **green** samples. At the second contact, an even larger dispersion occurs, as depicted by the **blue** samples.




Figure 4-3: Empirically measured trajectory of the first two contacts of a dice roll, overlaid with samples from the posterior probability distribution from the stochastic simulation. The samples are generated by the stochastic process used to propagate uncertainty.

Fig. 4-4 shows the predicted distribution of horizontal distances of the center of the dice normalized by the length of the dice from the simulation. The distribution broadens for each contact event, reflecting the increase in uncertainty due to contact.

The contact-free motion is governed by gravity and is deterministic, the first dynamic step is then from the initial pose to the end of the first contact, and the second dynamic step is from the end of the first contact to the end of the second contact. At contact the data-augmented contact model injects a measure of uncertainty into the belief over the states of the object that are then propagated along the trajectory. As expected, the propagated beliefs become more uncertain as contact occurs, as this is the main source of uncertainty in the simulator. In the following sections we first explain the details of the data-augmented contact models, then explain the nature of uncertainty and the method by which we propagate our belief.

In the results section, we show how this approach can be used to predict the outcome of planar dice throws.




Figure 4-4: Posterior distribution over horizontal distance of the center of mass of the dice normalized by the dice side length, from the stochastic simulator. Plots show initial configuration and the two post contact configurations.

4.4 Results

Residual learning

Fig. 4-5 shows the ℓ_2 norm error in predicted post contact velocity of the data-driven models compared to the Best Post Hoc model (best performance of the analytical models) and the IRB Bound model (best performance of any rigid-body contact model). The rigid models of either class are able to outperform the Best Post Hoc model, but only approach the IRB Bound model. These models are able to outperform their analytical counterparts, but because they are built using the same features and assumptions as the analytical models, they will not be able to exceed the IRB bound. In Class I, the \mathcal{X}_1 shows a marginally better performance over \mathcal{X}_2 , as it has a richer description of the impact event. The learned models with the extra explanatory power of wrench are able to outperform the IRB bound significantly and make far better predictions as they are not restricted by the same assumptions.

In Fig. 4-5(B), the *Data-Reinforced Rigid Par model* is the Anitescu-Potra model with model parameters as a function of the input state to the contact event. The performance of this model is close to data-augmentation with fixed parameters and shows that this paradigm is also viable.

Fig. 4-6 shows the expectation of the ℓ_2 -norm error in predicting the center of mass

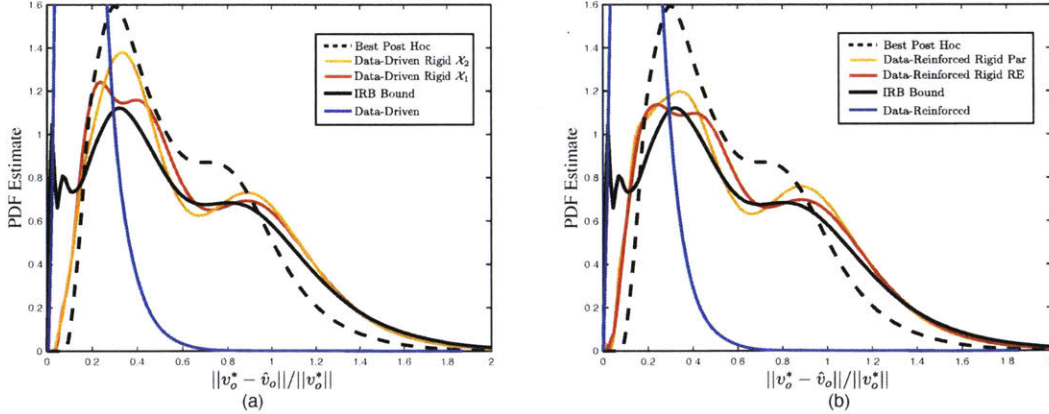


Figure 4-5: The estimated probability density function of the ℓ_2 norm error in predicted post contact velocity for 450 data samples used to train the models. a) Class I, b) Class II.

velocity as a function of the number of data used in the training, and here we show the best performing of each type of model for clarity. The data-driven and data-reinforced contact models are able to consistently and significantly outperform the IRB Bound due to their added expressibility, but require more data to converge (approximately 300 samples). The rigid learned models tend to converge to their optimal performance in approximately 100 samples, which is about double the number of samples required for the reliable identification of the parameters of the analytical models.

Uncertainty propagation

We validate the fidelity of the stochastic simulations on an empirically collected planar dice rolling data-set. Fig. 4-7 shows an example dice roll rendered from empirically collected data. The data is collected and processed as per the protocol used in [28].

We quantify the simulation fidelity by the likelihood of the ground truth final resting pose of the dice as predicted by the propagated belief. We compare the performance of the stochastic simulation to the analytic model. We further show how we can use the belief after an impact to grade the difficulty of predicting of dice rolls.

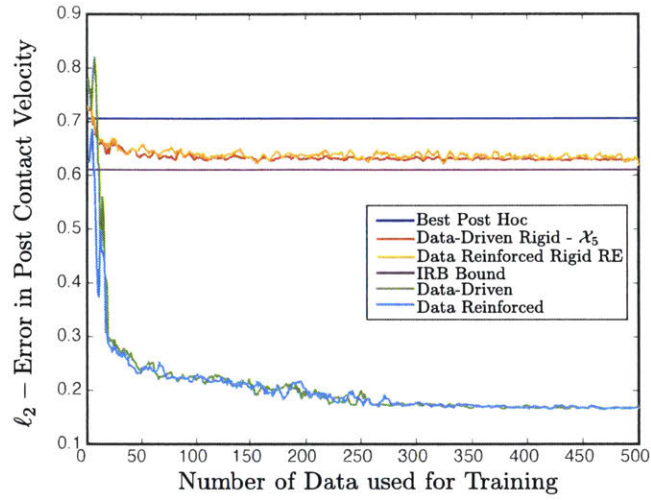


Figure 4-6: The ℓ_2 norm error in predicted post contact velocity for the models studied vs. the number of data samples used to train models.




Figure 4-7: An example dice roll rendered from empirically collected data. To generate the plot, we sample the trajectory measured from motion tracking uniformly and plot the trace of the outline of the block for each point.

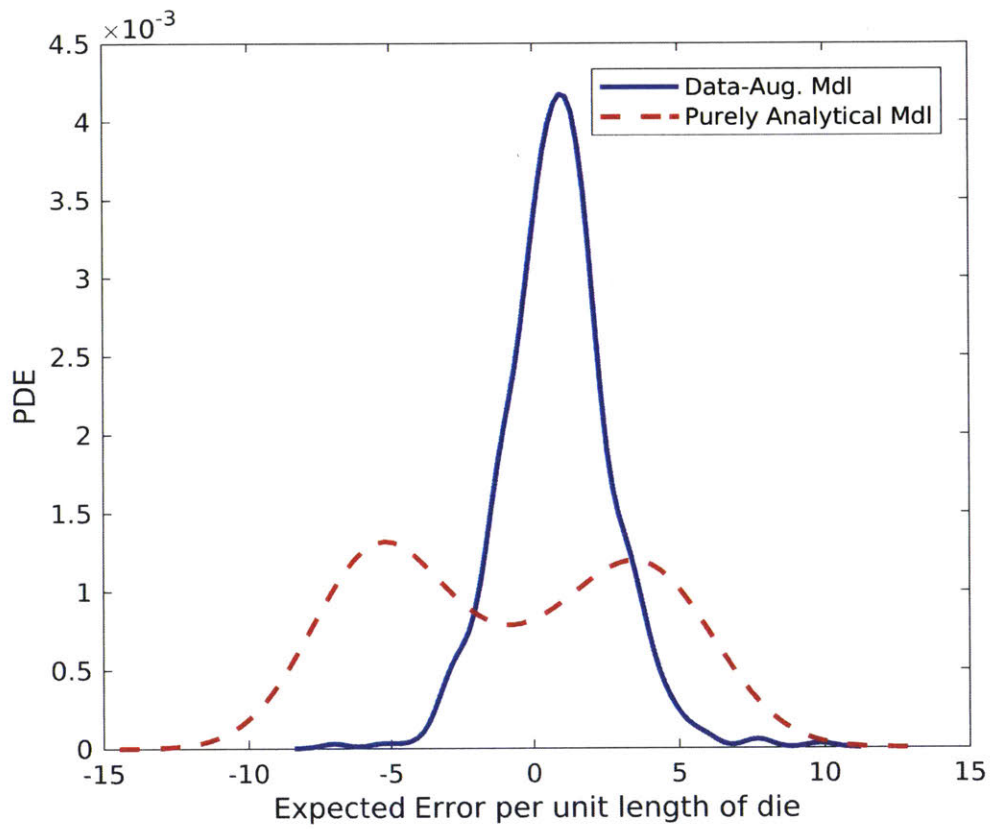


Figure 4-8: Error distributions over 750 drops normalized by the length of the side of the dice.

Fidelity

We evaluate the fidelity of the stochastic simulations by its accuracy in predicting the final pose of the dice as it comes to rest given an initial position and velocity prior to first contact. At rest, the dice height is equal to half the length of the side of the dice, so we quantify the error in horizontal position and orientation.

Fig. 4-8 shows the distribution of expected errors of the dice's final resting pose along the horizontal direction. The error is defined as the difference between measured and expected (from simulation) end position and can be positive (undershoot) or negative (overshoot). We normalize this error by the length of the side of the dice to give a notion of scale. A perfect model would have an exactly zero width peak at zero, this is equivalent to no error in predictions. The data-augmented model has a sharp unimodal distribution centered close to zero. This model exhibits lower variance when compared to the analytical model, meaning that it is more consistently accurate.

The expected error for the purely analytical and data-augmented models are respectively -0.97 and 0.83 , and the standard deviations are 4.44 and 1.83 respectively. Interestingly, the expected error across all trials for the two approaches is very similar (in absolute terms). The key difference is in the standard deviation and the spread of the errors (almost bi-modal in the analytical models case). The analytical model consistently either undershoots or overshoots the end configuration of the dice, and this seems to happen almost at random. This behavior leads to a close to zero average error over the entire data-set even though each individual throw is not accurately predicted. The over-shoot and under-shoot in this example are due to the sensitivity of the first contact to initial conditions, where the non-smoothness leads to a bifurcation in the end-pose of the dice.

The prediction accuracy of the data-augmented model is in line with the findings of [33], where for individual impacts these models were up to 3 times more accurate than their analytical counter-parts. The dice usually comes to rest after 4 bounces, though more or less bounces can be predicted from the simulations and these models are invoked each time.

The dice orientation takes one of 4 values at rest, we compare the most likely estimate of the simulation to the ground truth. If the two are in agreement, then we assign the drop a score of one, else we assign it a zero. The sum of the assigned scores over the total number of drops is the success rate of predicting dice roll outcomes. The success rate of the stochastic simulator was 43% and the analytical model 23%. We note here that random chance selection of the outcome has a probability of 25% to make a correct prediction. Interestingly, the analytical model seems to perform no better than random chance.

Utility of Belief Propagation

Belief propagation is a useful tool, consider Fig. 4-9 where we depict two throws taken from the test set. The impact on the left is harder to predict than the one on the right. This is because a clean corner impact with no immediate subsequent impacts is much easier to resolve than an almost plane contact with multiple impacts in quick succession. A small mistake in the latter can lead to vastly different predictions of end location of the dice.

Using the belief propagation approach we can find such examples by looking for impacts that have relatively narrower/wider posterior distributions over the set of outcomes. Fig. 4-10 shows the distribution over the resting pose along the horizontal axis for two such throws from the simulator. One throw exhibits a much larger variance when compared to the other (1.04 to 2.32 times the length of the side of the dice). We also overlay the expected outcome for each drop from the analytical model. We see good agreement between the prediction of the stochastic simulator and the analytical one for the throw on the right with lower variance and "easier" to predict. Instead, there is very little agreement between the data-augmented and analytical models for the difficult to predict throw, where only a small portion of the belief over outcomes agrees with the analytical outcome.




Figure 4-9: On the left is an example of a difficult to predict roll, where 2 temporally close consecutive impacts occur and a small error significantly impacts future contact events, whereas on the right a single clean impact occurs that can be predicted more reliably




Figure 4-10: Distribution over the possible horizontal resting positions of the dice for an easy to predict and difficult to predict roll normalized by the length of the side of the dice. The distribution is over-layed with the expected output of the analytical model.

4.5 Related Work

Recently, researchers have looked towards data-driven techniques to complement existing analytical models and/or learn dynamics directly from data. The work by [48, 1] are the closest to this thesis. In the former, a neural network is trained to provide inputs to an analytical model for planar pushing. In this framework, the output of the analytical model is used as the prediction; the neural network learns the best input parameters to maximize the performance of the analytical model. A benefit of this approach is that the model predictions are always feasible because of the analytical model, but the approach is deterministic, and relies on the expressiveness of the analytical model.

In [1], the authors proposed to learn a residual model for prediction of planar pushing. This approach is an extension of the work of this thesis to long horizon predictions using recursive neural network architectures. The key difference between [49] and [1] is that the former learns an input to an analytical model while the latter learns the residual error of the analytical model at the output.

In the planar pushing case, these models may be sufficiently expressive to span the full range of outcomes, but this is not always the case in other contact interactions as shown by [32]. Further, the models in [48] only make single-step predictions—an approach that may not work well for long-horizon predictions due to compounding errors at each time step.

In [95], the authors supplied a data-efficient approach to model the frictional interaction between an object and a support surface, by directly approximating the mapping between frictional wrench and slipping twist. Later, [96] extended the model to simulate parametric variability in planar pushing and grasping.

In [15], the authors showed how to design a neural network to predict rigid-body motions in a planar pushing scenario. In this study, as a robot pushes an object, the neural network differentiates between the object and the table. The neural network makes predictions by explicitly predicting $SE(3)$ transformations and jointly learning the full motion model and the observation model. This approach is still deterministic

and does not use any more physics knowledge.

4.6 Discussion

Contact models fulfill an important role in robotics and are used extensively in locomotion [54] and manipulation [19], yet the literature contains little work in empirical evaluation and comparison of the models prevalent in the field [84]. In this chapter, we show that empirical data plays an important role in highlighting the limitations of analytical models, and in allowing the construction of data-driven contact models that outperform their analytical counter-parts. We also demonstrate that by exploiting known structure and assumptions about physical interactions, we can build models that are more data-efficient. In particular, here we were able to reduce a 10-D space to 5-D. We further show significant improvements in model prediction when we allow models to deviate slightly from the requirements of rigid contact, i.e. using instantaneous wrench instead of linear impulse.

In this study we considered two classes of models, the purely data-driven (weak priors), and the data reinforced models (with strong priors). The challenge with using a purely data-driven model is the potentially poor or infeasible predictions made by these models before being trained on enough data. In contrast, augmenting analytical contact models with data allows for feasible predictions with adjustments made as more data is collected. With no data, the predictions rely on the prior generated by the underlying analytical models with no adjustment, which will be feasible but may be inaccurate. Our results suggest that these two formulations saturate at similar performances for this experimental setup, and the data reinforced models may be preferred due to the strong priors they provide.

The present study focuses on the impact of a single object, and learns a contact model suitable for the predictions of outcomes. Two interesting questions to ask are: how well does the model learned for this object work in making predictions for other objects of different shapes and physical properties? Is it possible to use this model as a prior and learn a correcting model with far fewer samples?

Chapter 5

Topic Models of Manipulation

5.1 Goal

In the previous chapters, we studied instances of manipulation models where a reasonable first principles models existed. In this chapter, we are interested in learning models for tasks without analytical representation. Further, we're interested in tasks that exhibit hybridness (due to parameter uncertainty or unobservability). For example, whether a block in a Jenga tower moves or not is dependent on unobservable frictional micro-interactions due to minute difference in block size and tower weight distribution or whether a door is locked or not depending on unobservable internal mechanisms of the lock.

Specifically, we present a two-step approach to learning a class of hierarchical models with application to physics modeling and robotic manipulation. These models allow temporal state abstraction learning and are amenable to probabilistic inference, controls, and planning.

To build intuition for the learned abstractions, we provide an analogy between these models and their natural language processing counter-parts.

We demonstrate the technique on a robotic system learning the mechanics of interacting with a Jenga tower. We show how the robot is able to sample-efficiently learn a model that it can then use for both high-level and low-level decision making.

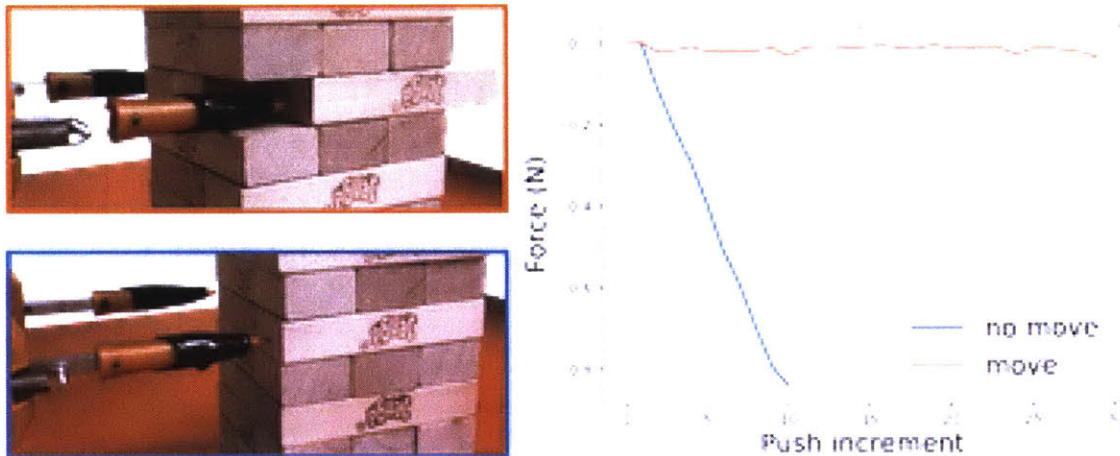


Figure 5-1: Model Learning in the Wild: A Robot learning the mechanics of interacting with a Jenga tower. Building useful temporal abstractions over these mechanics has powerful implications for prediction, planning, and controls. For example, inferring high-level block behaviors such as “free” or “stuck” early can simplify planning considerably. These behaviors are governed by micro-frictional interactions and unobservable geometric difference in blocks. Fortunately, tactile/force sensing provides rich contextual clues that help build abstractions that encode the effect of these mechanics.

5.2 Introduction

Building temporal state/action abstractions can significantly improve long-term decision making and predictive ability in agents [9]. For example, when opening a door, an agent may build useful abstractions over the mechanics of interacting with the door, i.e. whether it is locked or open. Inferring which of the two abstractions best describes the interaction fundamentally changes the way the agent acts (pulling the handle vs. reaching for a key).

Without these abstractions, the agent would need to reason through a long sequence of state observations such as door pose, handle pose, robot pose, and contact forces to make decisions. Typically, these abstractions have to be built in the wild, similar to how the robot in Fig. 5-1 would when learning the mechanics of interacting with the Jenga tower.

In this chapter, we present an approach to building abstractions through physical

interaction. Specifically, we present a two-step (top-down followed by bottom-up) approach to learning a class of hierarchical topic-model like representations of physics. These graphical models are inspired by topic models [40] and provide a mechanism to autonomously build temporal abstractions from state observations and actions. While in topic models the objective is finding latent structure in documents, here our objective is to find abstractions in the domain of physical states. Topics and abstractions are then analogs of latent structure in these two domains.

These models are potentially well suited to robotic manipulation since tactile measurements, much like words, provide rich contextual information that facilitates abstraction learning. Further, these models are well suited to inference, prediction, planning, and controls [13]. This chapter contributes:

- An extension and formalism for the modeling approach in [27].
- A grounding for the modeling approach through an analogy between the realization in the physical and Natural Language Processing (NLP) domains.
- A Demonstration of these models for a robotic system learning the mechanics of playing Jenga.

We begin the chapter by presenting the model and learning approach. We then provide the analogy between the physical and NLP domains. Next, we demonstrate the model in the Jenga experimental setup.

5.3 Formulation

Our objective is to build a hierarchical model that encodes temporal abstractions over physics and is amenable to inference, controls, and planning.

To illustrate, consider the example depicted in Fig. 5-2. The robot pusher interacts with the cube by pushing it. In (a), the block is very heavy and the pusher is unable to move it, while in (b) the block is light enough to move. Here, the state vector \mathbf{s}_t is composed of: i) the pose of the block (\mathbf{q}_t); ii) it's pusher location (\mathbf{p}_t);




Figure 5-2: Example: Here, a robot pusher is attempting to move the purple block. The block in (a) the block is very heavy and in (b) it is light enough to be displaced. How can the robot build useful abstractions such as “not movable” with physics predictions?

and iii) force measured at the pusher (\mathbf{f}_t). The goal is to construct a physics model that is capable of building useful abstractions such as “not moveable”.

To this end, we use the model depicted in Fig. 5-3. This model belongs to the class of hierarchical Bayesian models [46] and encodes information at 3 levels:

- The lowest level represents noisy observations \mathbf{z}_t of the states \mathbf{s}_t where subscript t denotes time. In Fig. 5-2, these correspond to the contact force, pose of the block, and the proprioception of the robotic pusher.
- The middle level represents a stochastic forward predictive physics model where the vertices connecting the sequence of states \mathbf{s}_t represent the conditional distribution $\mathbf{s}_{t+1} \sim \mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t, \mathbf{c}$. In Fig. 5-2, this corresponds to predicting the contact force, block pose, and robot proprioception given the robot pushing action.
- The highest level represents the abstractions made over the observed physics as latent variables \mathbf{c} . In Fig. 5-2, this could correspond to “movable” or “not movable”.

Without the abstractions, the model reduces to the standard Hidden Markov Model (HMM) where the middle (physics) and lower (observations) levels maintain their definitions. The model depicted in Fig. 5-3 bears close resemblance to Hidden




Figure 5-3: Plate Notation for the Hierarchical Physics Model: We can interpret the model as a filter of length D over N number of trajectories. K denotes the number of abstractions (unknown a priori). Details of model construction are provided in the text.

Markov Topic Models (HMTM) [3]. These models are commonly used in NLP to infer latent structure in corpora of documents. We provide a detailed analogy between the NLP variants and the physics model in Sec. 5.3.3. Here, we note two important distinctions between these two models:

1. the physics model is exogenous, i.e. that \mathbf{a}_t is an external input,
2. the states are sampled from continuous distributions as opposed to the countable, finite, and closed set of words in topic models.

The latter distinction has significant impact on parameter optimization and is the motivation for the proposed approach. The complex and multi-modal posterior distributions of the state transitions in the physics models, together with the interaction between categorical (abstractions) and continuous (physics) distributions complicates parameter optimization significantly.

5.3.1 Model realization

In this section, we provide details of the model realization in the physical domain. The model depicted in Fig. 5-3 defines a joint distribution over abstractions \mathbf{c} , states \mathbf{s}_t , state observations \mathbf{z}_t , and model parameters $(\boldsymbol{\theta}, \boldsymbol{\pi})$. Intuitively, the abstractions attempt to explain different physical regimes by building a segmentation of the state space with corresponding physics models. We can produce samples of state observations (observed variables) by conditioning on the other (latent) variables using the generative process:

$$\begin{aligned} \boldsymbol{\pi} | \boldsymbol{\alpha} &\sim \mathbf{GEM}(\boldsymbol{\alpha}) \\ \mathbf{c}_i | \boldsymbol{\pi} &\sim \mathbf{Multi}(\boldsymbol{\pi}) \\ \mathbf{s}_o | \mathbf{S}_0 &\sim \mathbf{S}_0(\cdot | \mathbf{v}_\theta) \\ \boldsymbol{\theta}_k | \boldsymbol{\Theta} &\sim \boldsymbol{\Theta}(\cdot | \mathbf{v}_\theta) \\ \mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t, \boldsymbol{\theta}_k, \mathbf{c}_i &\sim \mathbf{F}(\cdot | \boldsymbol{\theta}_{c_i}) \\ \mathbf{z}_t | \mathbf{s}_t &\sim \mathbf{Normal}(\mathbf{s}_t, \boldsymbol{\Sigma}_z) \end{aligned}$$

In this paper, we use the stick-breaking process [36] to construct the model with $\boldsymbol{\pi} \sim \mathbf{GEM}(\boldsymbol{\alpha})$ denoting the stick-breaking distribution over $\boldsymbol{\pi}$ [68]. $\boldsymbol{\pi}$ is the parameter of a multinomial distribution from which the abstractions are drawn and serves as a prior to the likelihood of each abstraction. This routine is implemented as:

$$\beta_k \sim \mathbf{Beta}(1, \alpha), \quad \pi_k = \beta_k \prod_{j=1}^{k-1} (1 - \beta_j), \quad \boldsymbol{\theta}_k | \boldsymbol{\Theta} \sim \boldsymbol{\Theta}(\cdot | \mathbf{v}_\theta), \quad \mathbf{G} = \sum_{k=1}^{\infty} \pi_k \delta_{\boldsymbol{\theta}_k}$$

This process provides a mechanism to sample from a variable number of abstraction with probability summing up to one. Intuitively, the abstractions represent a soft segmentation of the state-space \mathbf{S} into K finite measurable partitions S_1, \dots, S_K where K is also a random variable – stick-breaking allows sampling from randomly distributed K .

Here, $\mathbf{G} \sim \mathbf{DP}(\alpha, \boldsymbol{\Theta}(\cdot | \mathbf{v}_\theta))$ is the Dirichlet Process distributed parameters of the

model with concentration parameter α . This parameter sets a prior over the number of abstractions, where a larger value prefers models with higher abstraction counts. Theoretically, K can be infinite, but in practice we set the summation over π_k to a large positive integer $N \gg K$ and check the component values to evaluate truncation error. A rule of thumb for choosing N is $N \approx 10 \times \bar{K}$ where \bar{K} denotes the expected upper bound on the number of abstractions.

We denote the distributions over the initial state and weights of the state transition distribution with \mathbf{S}_0 and Θ respectively. These distributions are parameterized by the prior \mathbf{v}_θ . We assume \mathbf{S}_0 is normal distributed and an uninformative (by assigning a large prior to the variance of the distribution). We refrain from using the uniform distribution as an uninformative prior due to the complications introduced to momentum based optimization techniques from a lack of gradients.

The conditional transition distribution ($\mathbf{F}(\cdot|\theta_{c_k})$) relies on both the abstraction and the previous state. Physical systems are often complex and multi-modal, as such we choose expressive Bayesian neural networks (BNN) for $\mathbf{F}(\cdot|\theta_k)$ where θ_k denotes the stochastic weights of the k^{th} network. Let $\mathbf{v}_\theta := \{\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta\}$ and $\Theta = \mathbf{Normal}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$ then:

$$\begin{aligned} \boldsymbol{\mu}, \boldsymbol{\Sigma} &\sim \mathbf{Normal}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta) \\ \theta_k | \Theta &\sim \mathbf{Normal}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \\ \mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t, \theta_k &\sim \mathbf{BNN}(\mathbf{s}_t, \mathbf{a}_t | \theta_k) \end{aligned}$$

where we exploit the conditional independence in the graph for the state transition distribution. This formulation allows expressive multi-modal representations of physics. We assume the state observations \mathbf{z}_i are normally distributed with means centered at \mathbf{s}_t and variance $\boldsymbol{\Sigma}_z$, and infer the variance during learning.

To learn the model, we optimize over $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}_z, K\}$ to maximize the log-likelihood of the observations \mathbf{z}_t where K is the number of abstractions. Intuitively, the model is jointly inferring a segmentation of the state-space into while learning local physics models for each region. The optimization attempts to find this segmentation in the

joint domain of tactile and visual observations. Finding a meaningful segmentation is facilitated by the contextually rich but highly localized information provided through tactile sensing.

The challenge in learning this model is that state transitions are sampled from a continuous distribution $\mathbf{F}(\cdot|\boldsymbol{\theta}_{c_i})$ while the abstractions are drawn from a discrete distribution. In contrast, both distributions over abstractions and words of HMTM models are discrete, countable, and finite. While Gibbs sampling is an effective inference technique for the set of words in a corpus of documents, it does not perform effectively in sampling across discrete and continuous distributions with complex multi-modal posteriors. Furthermore, alternative variational inference schemes suffer from poor mixing and exploration because they must resort to Metropolis-Hastings Markov Chain Monte Carlo for the categorical distributions rather than using powerful gradient-based techniques such as No U-Turn Sampling. In the next subsection, we present a sub-optimal but numerically well-posed learning procedure for these physics models.

5.3.2 Model learning

Our approach to formulating a more numerically well-posed and sample-efficient parameter estimation optimization is to decouple the hierarchical abstraction learning from the transition learning. The intuition is that the abstraction that best describes a sequence of states and actions may not heavily depend on the details of the transitions between states (see Sec. 5.3.3).

In this approach, the agent first learns abstractions over the state observations (top-down learning [46]) while ignoring the temporal dependency between them (states are conditionally independent given the abstraction). Next, the details are filled in by learning the transitions (bottom-up learning), i.e. reintroducing the temporal conditional dependencies of the states. Our assumption is that abstract knowledge is sometimes available before more concrete knowledge is in place [45] and that hierarchical Bayesian models are an effective approach to capturing this form of knowledge [46]. The advantage of this break-down of learning the model is two numerically

well-posed and sample-efficient optimization programs for the model parameters.

During the top-down step, the agent decides the number of abstractions that describe the physical interaction ignoring temporal state dependence. This corresponds to removing the vertices connecting states in Fig. 5-3.

For the example in Fig. 5-2, this means that the robot-block interaction abstraction does not depend heavily on the order of the observations of state and action in a window of time. If the block does not move, we expect to see several large measurements of force but little to no block displacement irrespective of the order in which they appear.

This model can be written as:

$$\begin{aligned}
 \boldsymbol{\pi} | \boldsymbol{\alpha} &\sim \mathbf{GEM}(\boldsymbol{\alpha}) \\
 \mathbf{c}_i | \boldsymbol{\pi} &\sim \mathbf{Multi}(\boldsymbol{\pi}) \\
 \phi_k | \boldsymbol{\Phi} &\sim \boldsymbol{\Phi}(\cdot | \mathbf{v}_\phi) \\
 \mathbf{s}_t, \mathbf{a}_t | \phi_k, \mathbf{c}_i &\sim \mathbf{G}(\cdot | \phi_{c_i}) \\
 \mathbf{z}_t | \mathbf{s}_t &\sim \mathbf{Normal}(\mathbf{s}_t, \boldsymbol{\Sigma}_z)
 \end{aligned}$$

where the states and actions are distributed according to \mathbf{G} which is in turn distributed as a Dirichlet Process with a Multinomial base distribution. We have made a slight abuse of notation here in assuming that \mathbf{a}_t is chosen according to the abstraction. If \mathbf{a}_t is drawn from a joint distribution of abstractions and a policy, it can be incorporated in the model by extending \mathbf{G} to include the policy distribution accordingly.

We choose to use a Gaussian Mixture for the functional form of \mathbf{G} for three reasons: i) we assume the states and actions are drawn from a continuous support; ii) we assume each draw in a sequence of observations is independent and identically distributed (i.i.d.) (temporally uncorrelated); and iii) each sequence of states and actions is drawn i.i.d.. Let $\mathbf{G}(\cdot | \mathbf{c}_i) := \mathbf{Normal}(\boldsymbol{\mu}_{c_i}, \boldsymbol{\Sigma}_{c_i})$ where each abstraction \mathbf{c}_i represents a normal distribution with mean $\boldsymbol{\mu}_{c_i}$ and variance $\boldsymbol{\Sigma}_{c_i}$ (see Sec. 5.3.3 for details). In this step, the generative process provides a soft clustering of sequences of

observations using a variable dimension Gaussian Mixture. This allows the agent to localize and enumerate a set of k abstractions.

In the bottom-up step, we fix the abstractions and re-introduce the temporal dependence of the states. In the Fig. 5-2 example, we re-order the observations and learn the physics transitions assuming the abstraction. We now optimize over $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}_z\}$ to maximize the log-likelihood of the observations \mathbf{z}_t . Effectively, this step is learning a known number of mixtures with the BNNs as atomic models. One approach to aligning the mixture model with the previously acquired abstractions is to allow the mixture weights to depend on the states and actions explicitly [36].

This two-step optimization trades off numerical conditioning and sample-efficiency with optimality. The choice to attempt a one-shot optimization versus this two-step approach depends on the quality of the predictive performance of model. We recommend using the two step approach if a high-quality model cannot be found in one-shot. In the next subsection, we discuss the relationship between this model and the NLP variants. Further, we provide an analogy for the proposed optimization scheme.

5.3.3 Topic model analogy

Topic models are a class of generative probabilistic processes capable of finding latent semantic structure in a corpus of documents [40]. Document words serve as the observations for these models. The premise is that each document is a mixture of a small number of topics and that the probability of observing each word in the document is dependent on the associated document topics.

A simplistic interpretation of the semantic structure of a document (the topics) is a soft clustering of words that are “similar”. Similarity can be interpreted as the frequency of a given word/s conditioned on a set of latent variables that can be used to generate the observed documents. Optimization in these models is the joint inference of the number of topics and the degree to which any observed word belongs to any of the topics. Typically, words are assumed to be drawn from a prescribed lexicon (closed, discrete, and countable set of words) where a set of uninformative words such

as {"the", "and", ... } are omitted.

The model depicted in Fig. 5-3 can be interpreted as a HMTM [3] where z_t denote the observed words, the sequences of s_t denote sentences (an ordered set of words where the probability of occurrence of each word is conditioned on the previous one and the topic) and c_k denotes topic k . The inner and outer plates [11] are then the document and topic distributions, respectively. The transition weights θ_k parameterize the conditional probability of the occurrence of the next word given the current word. By erasing the vertices connecting s_t we arrive at the well-known Latent Dirichlet Allocation (LDA) process [12]. In the LDA model, only the occurrence of the words are considered, their sequentiality is ignored.

The analogy we make here is that state observations play the role of words, the sequence of states are sentences, and the abstractions are equivalent to topics. The key difference between the two models is the *support* (discrete vs. continuous) and *complexity* (low-dimensional vs. nonlinear, high-dimensional, multi-modal) of the mid-level transitions s_t . This difference poses a challenge in numerical optimization.

The intuition for the optimization approach is given a corpus of documents, we expect that both LDA and HMTM models produce similar semantic structure. This latent structure should not have a strong dependence on the order in which the words appear, rather just the probability of their occurrence. As [3] show, a HMTM model is capable of incorporating more fine-grained linguistic information and builds more refined but semantically similar topics to standard topic models. By ignoring the temporal dependence, we are effectively learning a standard topic model over physics, then augmenting this understanding with a HMTM-like structure for fine-grain details.

These models are particularly useful for manipulation due to tactile measurements. Similar to words, tactile sensory modalities provide rich semantic context. This semantic context is due to sparse, highly discriminative, and localized information provided through contact. For example, the existence of a contact force alone is highly discriminative in inferring contact formations which correspond to distinct semantics. The simplest semantic example is whether two bodies are in contact or not,

where the mechanics of two interactions are fundamentally different due to contact constraints.

In the next section, we demonstrate how these models can be used in the context of manipulation skills.

5.4 Experiments

In this study, we demonstrate the efficacy and sample-efficiency of a novel hierarchical learning approach to manipulation on the challenging game of Jenga. In Jenga, the objective is to extract blocks one by one from the tower and place them on top. Central to this game is the combination of sight and touch in inferring the latent state of blocks in the tower, i.e. semantics such as movable or immobile blocks.

The ability to infer whether the block is movable or not comes naturally to humans and influences their long-horizon decision making. This block behavior is governed by micro-interactions (minute differences in block sizes and frictional properties) that are unobservable through sight alone. By combining sight and touch, we show how a robot can build abstract notions of block behaviors through interactions that it can use for long-horizon planning and controls, inspired by our innate ability to do so.

In this section, we first describe the experimental setups (in the real world and in simulation). We then touch on the task specifications and metrics of evaluation.

5.4.1 Real-world setup

Fig. 5-4 (A) shows our experimental setup. The Jenga tower is placed in the workspace of a 6 d.o.f. robot manipulator (ABB IRB 120). The robot has access to an Intel RealSense D415 camera (RGB) and an ATI Gamma 6-axis Force-Torque sensor mounted at the wrist. The camera provides noisy estimations of 6 d.o.f. poses of the blocks using a specialized Mask-RCNN computer vision algorithm (see [27] for details). The force/torque sensor measures the forces applied to the robot due to the interaction.

We use the robots forward kinematics to estimate the pose of the gripper. To

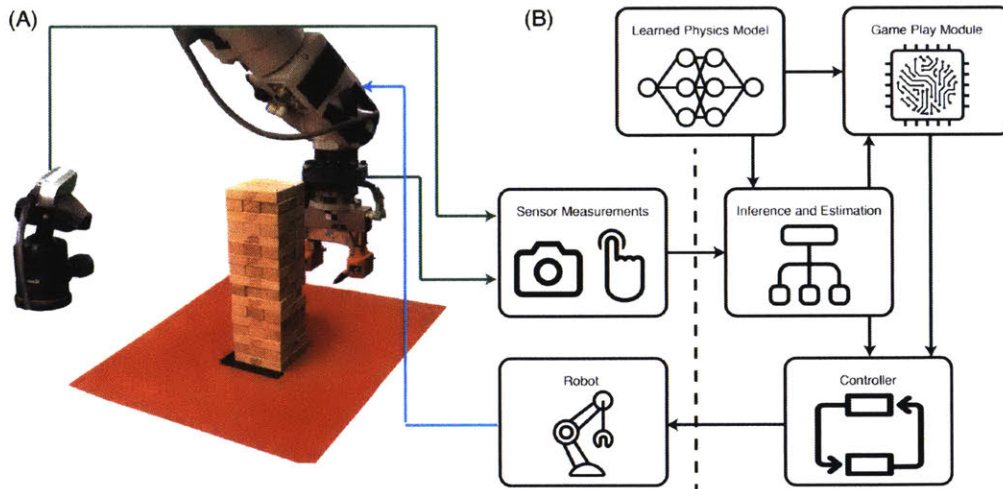


Figure 5-4: Robot Setup: Physical setup consisting of the robot, Jenga tower, Intel RealSense D415 camera, and ATI Gamma Force/Torque sensor (mounted at the wrist). (B) The machine intelligence architecture with the learned physics model.

estimate the pose of the fingertips, we compute the deflection of the fingers with respect to the gripper using the measured force applied to the finger and known compliance parameters. We use the experimental setup to demonstrate the fidelity of the proposed approach.

Fig. 5-4 (B) shows the software architecture we developed to run the system. The architecture is composed of:

- the learned physics model (the central focus of this chapter),
- a sensory module that processes RGB images and force measurements,
- an inference engine to estimate the ground-truth states,
- a finite-state machine (game-play module) that decides the next macro action of the robot,
- a control module that computes and communicates robot trajectories.

At each time-step, the robot calls its sensory module and runs inference on the measurements using its learned physics module. Based on the inferred states and

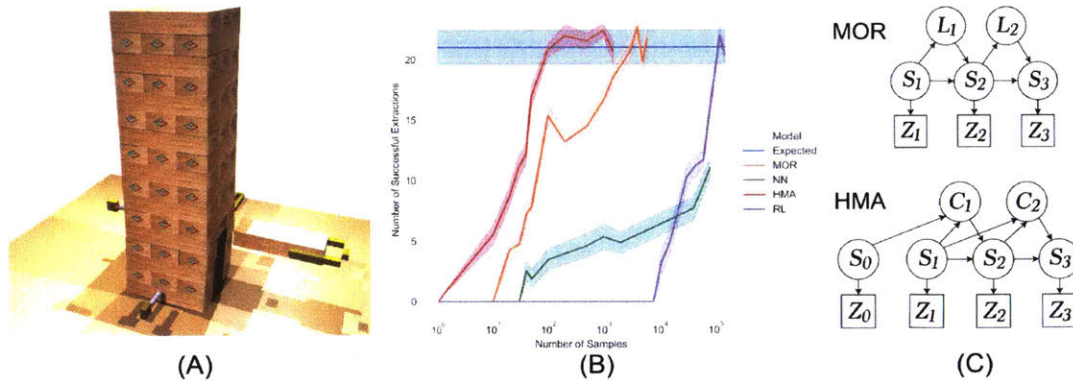


Figure 5-5: Jenga setup in simulation and the baseline comparisons: A) The simulation setup is designed to emulate the real-world implementation. B) The learning curve of the different approaches with confidence intervals evaluated over 10 attempts. C) A visual depiction of the structure of the mixture of regressions (MOR) and the proposed approach (HMA).

abstractions, the robot makes a decision to either push, retract, pick, or place the block. Each primitive has a low-level controller that executes it and can kick control back to the macro-decision maker given the new inferred states and abstractions.

5.4.2 Simulation setup

Fig. 5-5 (A) depicts our Jenga setup in the MuJuCo [84] simulation environment. In this environment, we simulate the robot with a point pusher for the push actions and a parallel jaw gripper for the extract and place actions.

The simulation environment shares the software architecture with the experimental setup. The simulated robot has access to the same set of forces and state measurements as the experimental counter-part.

The purpose of the simulation environment is to compare the performance of our proposed approach (Hierarchical Model Abstractions – HMA) to several standard baselines. In particular, data-collection for reinforcement learning is prohibitively sample-complex and is not practical for the real-world.

5.4.3 Task specifications and data collection

In this subsection, we specify the sensing modalities, actions, and rules by which the robot is allowed to play the game in both simulated and real environments:

- **Sensing:** The robot has access to its own pose, the pose of the blocks, and the forces applied to it at every time-step. The simulated robot observes these states directly, whereas the experimental robot has access to noisy estimates.
- **Action primitives:** The robot utilizes four “primitive” actions (push, retract, extract, place). Using the push primitive, the robot first selects a block and moves to a collision-free configuration in plane. The robot then selects a contact location and heading, and pushes for a distance of 1 mm and repeats. The action is considered complete if either the robot chooses to retract or a maximum distance of 45 mm is reached. The extract/place primitive searches for a collision-free grasp of the block and places it on top of the tower at a random unoccupied slot. The extract/place primitives are parametric and computed per call, as such they are not learned.
- **Base exploration policy:** The robot has access to a base exploration policy for data collection. This policy randomizes the push primitive by first selecting a block at random, then executing a sequence of randomized contact locations and headings.
- **Termination criteria:** A run, defined as an attempt at a new tower, is terminated when one of the following conditions is met: i) all blocks have been explored, ii) a block is dropped outside the tower or the tower has toppled.
- **Tower and robot specifications:** The simulated tower is composed of the same number and similar distribution of movable vs. immobile blocks as the real tower. This is due to slight perturbations to weight distribution resulting from small tolerances in the height of the blocks. The relative dimensions of the tower and the end-effector are consistent for both environments.

We define a sample of the push trajectory as the tuple $(\mathbf{z}_t, \mathbf{z}_{t+1}, \mathbf{a}_t, \mathbf{c}_t)$ (an incremental slice of the push trajectory). The exploration phase follows the robot’s nominal execution loop using the exploration policy with the purpose of collecting samples for model-training. Once data collection is complete, a model is trained and its fidelity is evaluated in a model predictive control framework per number of samples over an unseen set of test towers.

The robot’s nominal execution loop is to select a block at random and attempt the push primitive. During the push primitive, the robot either chooses push poses and headings or retracts. If the block is extracted beyond 3/4 of its length, the extract/place primitive is invoked. During a run, the nominal execution loop is continued until a termination criterion is met.

An important challenge is that movable vs immobile pieces are indistinguishable prior to contact and the robot needs to control the hybrid/multimodal interaction for effective extraction without causing damage to the tower. If the damage compounds then the tower loses integrity and termination criteria are met earlier. As such, this problem is a challenging example that motivates the need for abstract reasoning with a fusion of tactile and visual information with a rich representation of physics.

5.4.4 Evaluation metric

The complex, partially observable, and multi-modal mechanics of the blocks in Jenga poses a significant challenge to robots learning to play the game. These challenges are not exclusive to this game and exist in many manipulation skills such as assembly and tool use. As such, progress in effective learning of these skills is important and necessitates rich models and policies.

In this chapter, we evaluate the robot’s ability to play the game by counting the number of successful consecutive block extractions in randomly generated towers. This metric evaluates a model and/or policy’s ability to account for the complex and time-varying mechanics of the tower as the game progresses. This metric, and our study, emphasizes physics modeling and does not explicitly evaluate the adversarial nature of the game.

The robot can interact with a total of 45 blocks for a new tower. We found empirically that approximately 47% of pieces move in a real-world random tower and emulated this ratio in the simulation environment. Consequently, the robot can extract 21 blocks in expectation for an unperturbed tower of 3 by 18 (the last 3 layers are prohibited by the rules). This value provides a reasonable goal against which performance can be evaluated.

5.5 Results

In this section, we present quantitative results for our method and three competitive state-of-the-art alternative approaches in simulation on the evaluation metric proposed in the previous section. We then demonstrate the fidelity of our approach on a real-world implementation of the game. We end the section with an analysis of the physics and abstractions learned using the proposed approach.

5.5.1 Performance on evaluation metric

In this section, we answer the question “how much data does each model need to achieve the expected 21 block extractions in a single tower?”

5.5.2 Simulation

Fig. 5-5 (B) shows the number of blocks extracted using each model as a function of the number of samples used (simulation environment). For reference, a complete push is composed of 45 steps where a sample is measured at each step. A single sample takes approximately 2 seconds to collect experimentally and 0.3 seconds in simulation.

For the base-lines, we chose a feed-forward neural network as a representative non-hierarchical model-based approach, a mixture of linear regressors (MOR) model as a generic hierarchical model-based approach, and the proximal policy optimization (PPO) implementation of reinforcement learning as a model-free approach.

To train the models, random instances of the tower were produced (generating the approximate 47% movable blocks) and data was collected using the exploration policy. At certain intervals, models were trained and evaluated on a set of test towers. These towers were not used in training and were fixed for all models.

In this task, all models have access to the same set of states, actions, and model predictive controller (see [27] for further details on the controller). Fig. 2C depicts the schematics of our proposed approach and the MOR models. The MOR model makes use of latent variables l_t , and our HMA model uses abstractions denoted with c_t . The states and noisy observations are denoted by s_t and z_t , respectively.

The proposed approach reaches the expected number of successful consecutive extractions within 100 samples. The mixture of regressions model is next to achieve the maximum score, requiring an order of magnitude more samples. Interestingly, the feed-forward neural network saturated in performance, falling short of the expected maximum extractions. Upon closer inspection, we found that this model was unable to reliably predict the multi-modal interactions and either behaved too conservatively or too recklessly. The RL algorithm was the slowest in convergence.

Per run, all approaches were presented with new towers at random which explains, in part, why the RL algorithm takes a large number of samples to converge as the space of possible tower configurations is very large.

5.5.3 Real-world

Fig. 5-6 shows the robot’s performance prior to and post exploration. Here a successful push is one in which the robot is able to push a block to a desired end-goal without dropping it outside the tower or causing excessive damage. A successful extraction is one in which the robot is able to pull the block free after a push without damaging the tower and a successful placement is placing the block on top of the tower without damage.

The robot shows an appreciable improvement in block extraction, from 56% to 88%. The most noticeable gain is in side-block extraction, doubling the success rate from 42.6% to 78.3%. The results suggest that middle-block extraction is considerably

Block Position	Action	Exploration		Learned	
		Attempts	Successes	Attempts	Successes
All	Push	403	172 (42.7%)	203	93 (45.8%)
	Extract	172	97 (56.4%)	93	82 (88.2%)
	Place	97	85 (87.6%)	82	72 (87.8%)
Side	Push	288	122 (42.4%)	133	69 (51.9%)
	Extract	122	52 (42.6%)	69	54 (78.3%)
	Place	52	44 (84.6%)	54	49 (90.7%)
Middle	Push	115	50 (43.5%)	70	33 (47.1%)
	Extract	50	45 (90.0%)	33	28 (84.8%)
	Place	45	41 (91.1%)	28	23 (82.1%)

Figure 5-6: Summary statistics for Exploration and Learned Physics. Table comparing the performance of the robot using the exploration strategy and the learned model.

easier than the side blocks due to its constrained motion and the favorable weight distribution of the tower. The robot is able to displace 42.7% of the blocks, close to the empirical average of a random tower.

The failure criteria is expanded to include tower rotation or displacements exceeding 15 degrees and 10 mms respectively. This criterion is imposed by the poor predictions made in the vision systems beyond these values. The exploration strategy is also modified to include a hand-coded supervisory algorithm that attempts to mitigate damage to the tower using measurements of poses and forces but is tuned to allow mistakes.

The two main failure modes for the extraction of blocks are: i) excessive forces applied to the tower (characteristic of failing to identify block behaviors), and ii) poorly controlled extraction of blocks (characteristic of poor predictive ability). The first failure mode often results in either a large tower perturbation such that the vision system is no longer reliable or the tower collapses. The second failure mode

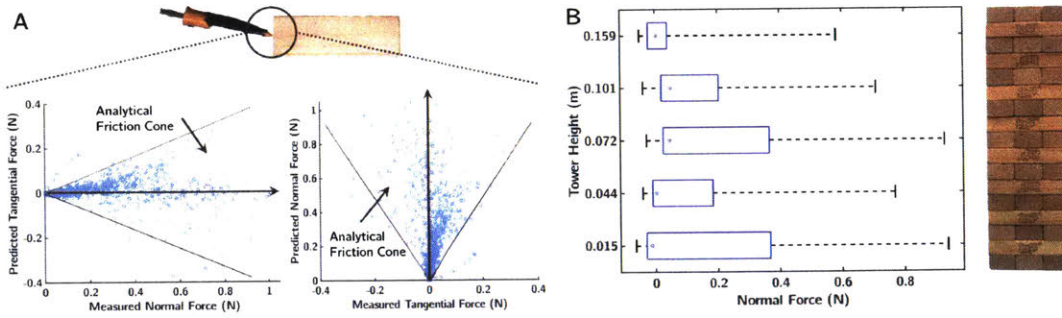


Figure 5-7: Learned Intuitive Physics: (A) Overlay of the analytical friction cone and predicted forces given the current measurements. The friction coefficient between the finger material (PLA) and wood is between 0.35 and 0.5, here we use 0.42 as an approximation. (B) The normal force applied to the tower as a function of the height of the tower.

often leads to either blocks dropping outside the tower or ending in configurations that are either difficult to grasp or occluded to the camera.

5.5.4 Model learning

In this section, we pay closer attention to the models learned for the experimental case. Fig. 5-7 depicts two examples of the physics learned by these models. Fig. 5-7 (A) depicts the analytical friction cone between the fingertip and block overlaid with the predicted normal and tangential forces given their current measured values. This cone is computed assuming Coulomb friction and rigid point contact between the fingertip and block. Under these assumptions, any transferable force between the finger and block must lie on the boundary or in the interior of these cones. The predictions of the models are in good agreement with the cone, implying that it has learned some latent representation of friction. We note that the friction cone is invariant to the abstractions and the predictions of the models reflect this fact well, implying coherency across models and abstractions.

Fig. 5-7 (B) shows the resistive force of blocks as a function of the tower height, measured vertically from the base of the tower.. The model is able to capture the intuitive tendency of block to resistance to decrease with tower height (due to the de-

crease in effective remaining weight on top of each block). The abstractions facilitate this by partitioning the state-space between macro-block behaviors and efficiently differentiating between movable vs immobile blocks.

Inference, controls, and planning

The abstractions offer coarse but intuitive and interpretable explanations of physics modes. As such, knowledge of the particular mode can be effectively leveraged for controls and planning. As a first step, the robot must infer the abstraction from noisy sensor measurements. Once learned, our representation provides a joint distribution over the states and abstractions, i.e. a generative probabilistic model of physics.

We use Markov Chain Monte Carlo sampling with Hamiltonian dynamics to approximate the posterior distribution over states and abstractions given noisy measurements. Fig. 5 shows two examples of the inferences of block types made by the robot during a push trajectory. The force and displacement plots show the prediction of the state for each abstraction. The abstraction probabilities are directly related to how well they are able to explain the observations.

In the case of a block that does not move, Fig. 5-8 (A), as there are no changes initially in force or configuration, the robot assigns the most probability mass to “no block” and some probability to “easy move” as the force measurements match fairly well. As the force climbs but configuration does not change, the probability of “no move” and “hard move” increase. The “hard move” abstraction does not predict the continued climb in force with no displacement and loses likelihood in the next few steps.

Fig. 5-8 (B) shows a block that moves with very little resistance. During the initial phase, due to the lack of force and displacement, the robot assigns the most probability to “no block”. As the push continues, displacements increase while forces do not, shifting the belief to “easy move” with very low probabilities for the other abstractions. Qualitatively, the inference is mostly driven by force measurements during the initial phase of a push because of the relatively small changes in configuration.

Critically, the perception algorithm has the highest uncertainty in the beginning

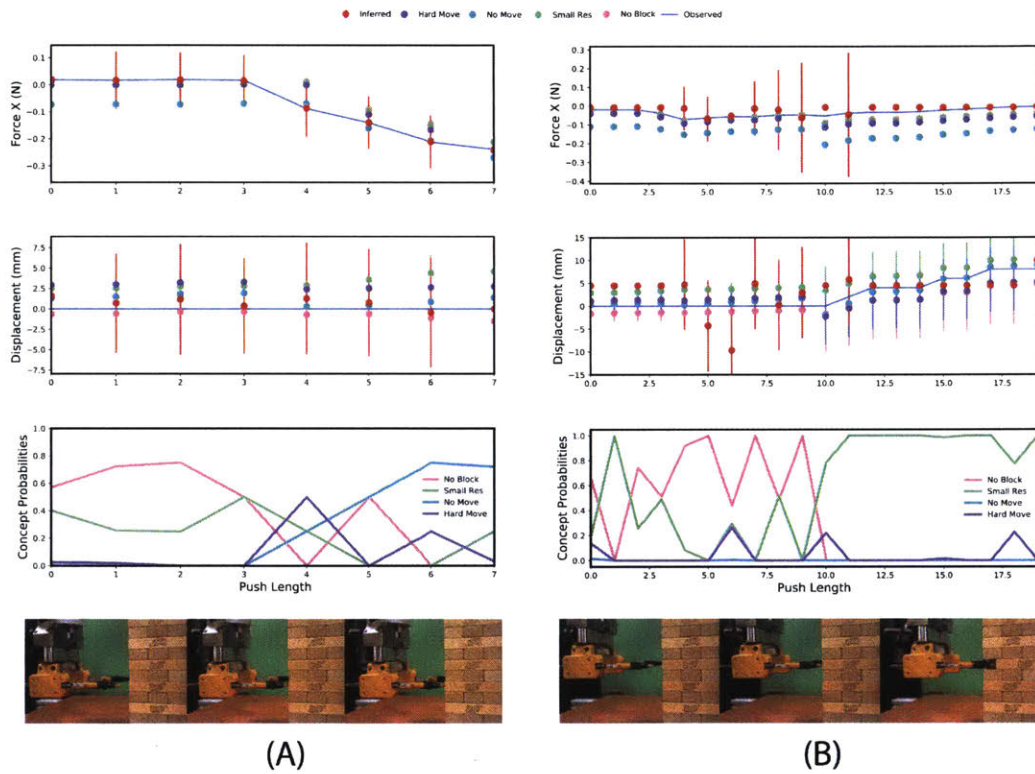


Figure 5-8: Inference using the learned representation: Evolution of the beliefs of the robot as it interacts with the tower. (A) For a block that is stuck. (B) For the block that moves easily.

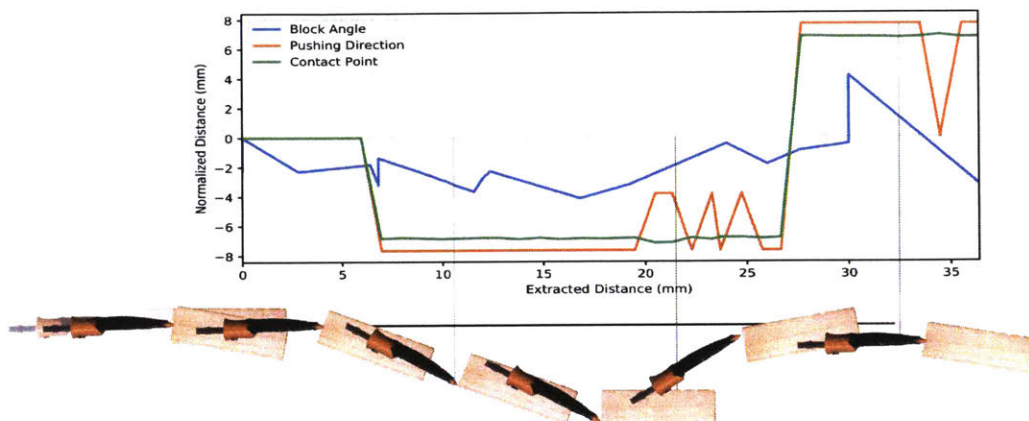


Figure 5-9: Controlled block pushing: The robot selects the point on the block and the appropriate angle to push with such that it realigns the block with the goal configuration. Here, the block is beginning to rotate counter clockwise and is starting to move out of the tower. The robot selects a point close to the edge of the block and pushes it back in towards to tower center. We convert angles to normalized distances by scaling with the radius of gyration of the block, 0.023 meters. We have exaggerated the block translation to illustrate the fine grain details of motion.

of the push as the mask computed for the block is the least informative and many hypotheses explain the current pose. Configurations play a more decisive role further along the push as they become more certain and discriminative between abstractions. This observation highlights the importance of inference and modeling in the joint domain of visual and tactile modalities.

To use the abstractions in a control or decision-making framework, we first need an association between desirable abstractions and actions. To incorporate this information, we add a cost term to the MPC controller, where actions resulting in increased probability of "no move" are penalized. This approach is very similar to standard reward shaping for the model-free reinforcement learning frameworks. A desired policy is one that leads to minimal damage to a tower. We impose this by assigning a large cost to tower perturbations which is effectively an implicit encoding of our human intuition.

Fig. 5-9 shows an example trajectory for the block and the controller actions.

Here, the robot attempts to push the block to a desired end-configuration using its MPC controller. Using the inferred abstraction and states, the robot forward predicts the costs of sequences of actions over a time-horizon, then executes the first action from the sequence incurring the smallest cost and repeats. The fidelity of the representation in forward prediction coupled with knowledge of the abstraction provides effective fine-grain control of block motions. This in turn significantly mitigates dropping blocks outside the tower or toppling, allowing runs to continue for longer.

5.6 Related Work

The main objective of this chapter is to propose an approach to learning temporal abstractions over physics. The ideal goal of this line of research is to use these autonomously acquired abstractions for long-horizon decision making (whether learned or planned).

Planning at multiple levels of temporal abstraction has been a long-standing challenge in the RL and AI [82]. Attempting to learn long-term policies without temporal abstractions often leads to prohibitively sample-complex scenarios and limited re-usability of acquired policies.

A notable sum of seminal work in hierarchical reinforcement learning addresses the model-free scenario. Works such as Options [82], Feudal Learning [20], Hierarchical Abstract Machines [67], and MAX-Q [22] have focused on carefully hand-crafted hierarchical structures with abstractions implicitly provided by experts. In these models, the abstractions are carefully hand-crafted. The hierarchical structures used in these frameworks typically induce a Semi-Markov Decision Process (SMDP) over the MDP studied in vanilla RL.

More recently, researchers have moved towards autonomous abstraction learning. The literature contains a wealth of extensions to the traditional frameworks for greater autonomy. The works of [63] and [23] are some of the closest work to this chapter. Both works propose to learn generative hierarchical representations. In particular, the latter uses conditional restricted Boltzmann machines (CRMBs) to build abstractions.

These un-directed graphical models are the cousins of the directed graphical models discussed in this paper.

By learning abstractions over physics, we can provide model-free HRL approaches with an autonomous mechanism to infer “initiation sets”, “termination conditions”, and physics-based factorization of the state-space. This alleviates the challenges in hand-crafting substructure and local policy learning. The abstractions learned from physics can also serve as inputs to the value functions used in approaches such as [82].

On the planning side, a hierarchical factorization of the state-space may introduces compositionality in plans. [43] investigates learning composable models and a formal approach to planning using “backchaining” – a process whereby a plan is computed backwards from the goal by sequencing subroutines with conducive preconditions.

[88] investigates the likelihood of success of a sequential plan given precondition parameters of subroutines. Planning approaches are more robust to changes in goals and deal better with previously unseen task constraints. Similarly, [85] investigated a hierarchical planning approach to long-horizon manipulation skills. The central idea in both approaches is a definition of an macro-action (such as pushing or pouring). Our objective is to help these approaches in less structured environments build macro-actions using abstractions acquired from exploration and model-learning.

From a meta-learning perspective, our approach looks to organize instances of the physical world into categories that it can use to plan or learn actions. [34] demonstrates a model-free approach that leverages gradient-descent to adapt the pre-trained parameters of a general-purpose regressor to new tasks. Physical systems tend to have latent structure that can be exploited to build simpler policies that can more efficiently adapt to new tasks.

Alet et al. [2] demonstrate a graph neural network approach that learns compositional physics structure that can be re-arranged to generalize to novel configurations. The abstractions in our physics model serve as indices that weight the predictions of the atomic models. By combining representations such as [2] together with hierarchical learning, we may be able to build more expressive models with better

generalization properties by introducing additional mechanisms for compositionality.

From a latent structure learning perspective, [16] explore latent structure learning from task demonstration for trajectory generation in the for collision free robot motion. Further, in [37] abstractions such as soft or hard are used in meal preparation, where physical attributes of food items are learned then used to prepare a salad. Learning fine-grained physical manipulation skills is still an important open problem in this domain.

From a Bayesian hierarchical model-learning perspective, the work of [40] and [12] is a key inspiration, where we have adapted topic models to manipulation. These works are discussed in more detail in the previous sections.

5.7 Discussion

In this chapter, we took a top-down bottom-up approach to manipulation skill learning. This approach to top-down learning is inspired by “concept learning” as studied in cognitive science. A “concept” is a type of abstraction composed of a typical feature-set. We naturally categorize objects and ideas into concepts, for example doors that are locked or open. Concepts are not measurable, rather they provide useful abstractions for categorization without the need to specify fine-grained details of instances.

When playing Jenga, we intuitively categorize blocks based on their general behaviors. For example, we may refer to blocks that do not move and exhibit large resistance (the typical feature-set) as “stuck”. From experience we learn that this type of block does not help us make progress in the game. We infer connections between concepts and our objectives but these abstractions are made independent of the task, rather they are a latent property of the physics. These coarse abstractions organize our reasoning and facilitate decision-making.

From a robotics perspective, there are two important benefits to learning a few general concepts then learning fine-grained physics models. First, by virtue of implicitly factoring the state-space and sequentially learning physics using simpler local

models, we can significantly increase the sample-efficiency of learning (Fig. 5-5). Second, concepts may capture interpretable modalities in physics that can be used in controls and planning. In this study, we take a nongoal-directed perspective, first determining these abstractions then learning models. This approach is related to [28] where a goal-directed latent space representation of the task is inferred in the context of learning from demonstration.

Certain manipulation tasks, such as meal preparation, may benefit from other abstractions such as soft or hard. In [27], a latent representation of physical attributes of food items is learned that are then used to prepare a salad.

Mixture models are an alternative approach to learning latent states and predictive models. These models are particularly effective in optimizing for predictive ability by simultaneously computing the number of mixtures and model parameters to maximize a likelihood score. A challenge in mixture models is making effective use of the latent variables for model-based controls and planning, in particular when the number of mixtures is large. For example, in the simulation environment the number of mixtures ranges from 2 to 11 (determined using 5-fold cross-validation) while the underlying latent structure does not provide additional insight that we found useful to controls or decision making. An additional technical challenge in mixture model learning is parameter estimation. Expectation-maximization is commonly used for estimation but can be sensitive to output variance, initial conditions, and local-minima.

Sample-based variational inference approaches do not scale well with data size and model dimensions, in particular due to the simplex constraint on mixture weights and poor mixing of the samples due to the complex topology of the posterior distribution. In both instances, the potentially large number of mixtures further increases the parameter space, requiring more data to learn and partially explains the results of Fig. 5-5 (B).

Mixture model learning is closely related to idea of dynamical mode learning. Dynamic modes are partitions in state-space for which transitions are smooth. Inferring dynamic modes is an important challenge in robotics, however, the level of detail provided by such a description may not be necessary for successful task execution. For

example, for a moving block in the Jenga tower there are at least 8 distinct contact formations, 4 of which are unstable, and each has its own dynamics. The subtle differences in the modes renders them difficult to infer and significant expert knowledge is required to render them useful. Learning coarse abstractions can provide sufficiently descriptive representation of the interactions.

An alternative approach to model-based manipulation skill learning is policy learning. Recent trends in reinforcement learning for manipulation have mostly focused on learning in the visual domain, leveraging advances in deep reinforcement learning and computer vision algorithms (19, 20). These algorithms are used to produce autonomously acquired manipulation policies that map from the pixel domain of images to robot actions. Real-world policies learned in the image domain have been shown for planar pushing, bin picking, and insertion [21–25].

These approaches work well for tasks in which the visual data-stream provides sufficient information and data collection can be automated effectively. There are three important challenges in learning policies for manipulation. First, many contact rich manipulation skills are difficult to automate for large-scale data collection. Jenga is a prime example where tower resets are time-intensive; therefore, sample-efficiency is an important concern. Second, Sim-to-Real transfer of policies remains challenging as most simulators use computationally efficient but inaccurate models of frictional interaction [26].

In the case of Jenga, the block motions rely on very fine micro-interactions and pressure distribution variations that are not observable or measurable in practice. Third, tactile information is often intermittent, i.e. making and breaking contact over a short duration. Consequently, the effective integration of tactile information together with the persistent visual stream is challenging. These challenges are prevalent in many manipulation tasks such as small parts assembly, warehouse logistics, and disaster response. In such tasks, fine level reasoning for contact is critical and mistakes often incur large costs.

The interplay between tactile and visual information is central manipulation. These two modalities act as complements and provide information in the absence of

one-another, for example, visual occlusions are a common part of Jenga, where blocks are occluded during manipulation but the tactile modality continues to provide information. Further, tactile feedback can provide high-resolution local information to complement the global but coarse information from vision.

In the experimental setting, our perception system is susceptible to errors in pose estimation during the initial phase of the push as the total visible mask of the block is small and several hypotheses may be equally likely. The force-stream and proprioception, together with their temporal variation increase the likelihood of plausible hypotheses for poses despite the occlusion of the end-effector due to the tower.

Building coarse abstractions in the joint domain of tactile and visual feedback is central to our approach. These abstractions are partially driven by the temporal change in these signals, motivated by time-varying multi-modal mechanics. The effective use of both modalities together with the abstractions leads to significant sample-efficiency, allowing the experimental robot to learn tower physics entirely from real-world experience. Furthermore, the abstractions are effectively used for controls and planning, significantly improving the performance of the system. In particular, for Jenga our controller is able to explicitly use the relation between inferred stuck pieces and tower perturbation to mitigate damage.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Closing Thoughts

In this thesis, we covered a wide spectrum of topics. Our discussion took us from analytical model-based approaches to fully data-driven ones. In this last chapter, I wanted to take the opportunity to discuss a few things I have come to believe over the last five years as of the time of writing this thesis. These thoughts have shaped much of the underlying themes of this thesis and serve two purposes: i) a grounding of the approaches we explored, and ii) a discussion on future work.

6.1 Why model-based or compositional approaches?

There's a large body of work addressing model-free approaches using reinforcement learning techniques. These approaches use either pixels from RGB cameras, raw motor encoding, or some combination thereof. These methods look to compute a policy for given state observations to maximize a reward function that serves as a proxy to a goal. Bypassing the need for models and representations seems to be an attractive idea. If you can make decisions in one shot from raw sensory data, why not go for it?

These approaches are a great idea in environments where the state vector is fixed and fully observable and experiments are easily repeated. Chess is one good example. While extremely complex, the state vector of chess is fixed, the pose of every piece is always perfectly measurable and the outcome of every action known. Chess is also

cheap to simulate, and the simulation matches the real-world perfectly.

Why is manipulation any different? Let's consider a simple daily task of opening a door; a skill we would expect our robots to master to be useful. How many types of doors can we imagine? Hinged doors, sliding doors, vertical doors, hatches, etc.. How many different handles can we imagine? Rotational, pivoting, sliding, etc.. How many different ways can a door open? How many states (unobservable prior to interaction) can a door be in? Locked, open, not a door, etc.. How many different colors are doors? How many different sizes? Different handle locations?

You get the idea, there is such a vast variety in one of the most standard manipulation tasks we can imagine, how would we simulate and test all these possibilities at scale? How many forests would we burn down for this one task? How many forests would we need to scale to basic manipulation skills? Contrast this to a human who reasons explicitly about doors as a concept that itself is composed of concepts of different kinds of doors and behaviors. One natural way to emulate this reasoning is to use compositional models and policies that adapt to specific instances of a task. Hierarchical Bayesian model learning and meta-learning are examples of approach that aim at exactly this by explicitly baking in compositional/hybrid structure. Chapter 4 of this thesis demonstrates one example of these approaches.

We are still very far away from a robotic system that can autonomously acquire such abstractions in manipulation. The hope is that by providing these systems with innate knowledge of physics and the ability to acquire abstractions, we can build systems that are more autonomous in the wild.

To this end, learning abstractions should not only be over physics (states/actions) but also over relations. For example, what does it mean for two bodies to be rigidly attached? We write this down as a kinematic constraint. Can robotic systems build similar abstractions (without enumerating all the possible rigid body constraints) and does planning benefit from explicitly reasoning over these?

We refer to the idea of learning policies or planning that is not only conditioned on the states, but also the constraints or abstractions. Rather than hand-coding or enumerating them, having a set of them learned.

6.2 On simulation to real-world transfer (sim2real)

Solving manipulation is not just about domain randomization in simulation followed by sim2real. There are a few reasons why.

Currently, there are few if any reliable simulators that can produce tactile observations with any reasonable fidelity. Force/torque sensing from simulators is not reliable, and tactile sensing from sensors such as Gelsight (tactile imprints of objects they touch) is not available. This has incentivized some researchers to claim that tactile sensing is not necessary for dextrous manipulation. What these researchers fail to mention is that without this modality, they would not be able to reach into their pockets and retrieve their keys or notice that they are touching a hot plate.

Some problems are too difficult to simulate reasonably without significant engineering knowledge about the problem itself. Take Jenga, if we did not have the insight that the micro-interactions between blocks was due to minute differences in block size (imperceptible to the human eye), how would we go about simulating the game? Engineering simulation setups for complex manipulation skills is non-trivial. Solving this problem may often lead to solutions intuited by the engineer that are better than those acquired by the agent. Rather than the agent learning the task, it is the engineer who then coaches the agent to learn.

It is a well documented phenomenon that agents who learn in simulation environments often evolve strategies that exploit the simulators. There are instances of agents that are supposed to learn to walk, but rather learn that by forcing their feet into the ground, they can gain significant vertical impulses that they use to reach their goals by leaping. Automating a suite of manipulation learning simulators would require careful pruning of the many possible corners cases that arise. Given the vast number of possible manipulation skills, how do we address this?

Our proxy to learning policies is a reward function we'd like our agents to optimize. Engineering this reward function maybe almost as difficult as solving the task itself. For example, what would a reward function for tightening a lid on a coffee mug look like? Given the limitations in simulators that we have, this problem is further

aggravated.

In this thesis, we saw multiple instances in which the simulator would not match up well with the empirical data. Evaluating the fidelity of our simulators for robotic manipulation over a distribution of tasks is still an open problem. Such studies would begin to answer some of the key concerns in the sim2real transfer debate.

6.3 State-estimation and perception are hard

Part of our innate ability, as humans, to reason over objects and relations is our ability to infer their existence and states in a given scene. Despite the major progresses in computer vision, object segmentation and pose estimation in unstructured environments is still quite challenging. Pose estimation accuracy is important in fine manipulation skills. For instance, a mistake of just a few millimeters is the difference between a collapsed or stable tower in Jenga, or when levering an object it is the difference between missing the support or not. This estimation often relies on the a priori knowledge of the existence of specific objects/bodies in the scene given by the engineer.

Currently, our best bet in reaching this level of accuracy is to tailor highly specialized perception systems that are often over-fit to instances of tasks. This engineering effort is significant and often a requirement for the task, not the objective of the research on that task. Progress in state-estimation algorithms that leverage RGB-D sensory modalities together with touch and physics filters is an important step towards autonomous robotic manipulation. We, as humans, are successful at eliminating implausible hypotheses such as poses of objects that lead to penetration. While some work exists in the robotics and computer vision literature, it is still in its infancy and there's much room for progress. Critically, demonstrating algorithms that work robustly in the real-world is likely to progress manipulation further than showing results on specific benchmarks (that often are not correlated to practical physical tasks).

Underlying each chapter of this thesis was a significant amount of engineering for perception. The earlier chapters relied heavily on AprilTags (a type of QR code tag)

and Vicon motion tracking while the latter chapters relied on Mask RCNNs (a type of segmentation to pose estimation algorithm) that were over-fit to the specific task with significant engineering effort. While necessary and important, this engineering effort is often not the goal of research in manipulation skills. Some good tools such as DOPE [86] and DART [75] have been developed, we hope to see more in work in this area.

6.4 The action-perception loop

In a nutshell, all our autonomous systems are looking to take actions, evaluate how the world changes, and update their internal belief, then repeat the cycle until they acquire enough information to achieve their objectives – the action-perception loop.

My belief is that the lower-dimensional encodings available from physics and coarse abstractions are an effective way of encoding and filtering information. In the first few chapters of this thesis, we assumed that the perception model (our observations of states and object relations) was correct (an oracle). For example, in the planar impact task, we knew there was an object, and it made frictional interaction. We tracked the object and assumed the state-observations were sufficient for identification.

The the final chapter, we relaxed this assumption a little by not assuming anything about the interaction between blocks, rather learning a proxy to it through abstraction learning.

The most capable robotic systems will be those able to refine their perception models and their action/physics models in situ. Belief space planning [70, 69] in domains where the number of abstractions may increase over time is an interesting direction. For instances, the robot playing Jenga may in the beginning believe the world is composed of only a rigid tower, then it may refine itself model to represent the tower as pieces, then as pieces with properties. How do we plan for or learn in the domain of variable beliefs in the structure of the environment without enumerating every-single possibility (assuming that were possible)? What the the optimal exploration actions in these scenarios?

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

- [1] Anurag Ajay, Maria Bauza, Jiajun Wu, Nima Fazeli, Joshua B Tenenbaum, Alberto Rodriguez, and Leslie P Kaelbling. Combining physical simulators and object-based networks for control. *arXiv preprint arXiv:1904.06580*, 2019.
- [2] Ferran Alet, Tomás Lozano-Pérez, and Leslie P Kaelbling. Modular meta-learning. *arXiv preprint arXiv:1806.10166*, 2018.
- [3] Mark Andrews and Gabriella Vigliocco. The hidden markov topic model: A probabilistic model of semantic representation. *Topics in Cognitive Science*, 2(1):101–113, 2010.
- [4] Miha Anitescu and Florian A. Potra. Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Non-linear Dynamics*, 14:231–247, 1997.
- [5] Mihai Anitescu and Florian A. Potra. A time-stepping method for stiff multibody dynamics with contact and friction. *International Journal for Numerical Methods in Engineering*, 55(7):753–784, November 2002.
- [6] Christopher G. Atkeson. State estimation of a walking humanoid robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3693–3699, 2012.
- [7] Ko Ayusawa, Gentiane Venture, and Yoshihiko Nakamura. Identifiability and identification of inertial parameters using the underactuated base-link dynamics for legged multibody systems. *International Journal of Robotic Research*, 33(3):446–468, 2014.
- [8] SV Barai and PC Pandey. Vibration Signature Analysis Using Artificial Neural Networks. *Journal of Computing in Civil Engineering*, 9(4):259–265, 1995.
- [9] Andrew G Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1-2):41–77, 2003.
- [10] Maria Bauza and Alberto Rodriguez. A probabilistic data-driven model for planar pushing. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3008–3015, 2017.
- [11] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.

- [12] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- [13] Matthew Botvinick and Marc Toussaint. Planning as inference. *Trends in cognitive sciences*, 16(10):485–488, 2012.
- [14] B Brogliato, AA ten Dam, L Paoli, F Génot, and M Abadie. Numerical simulation of finite dimensional multibody nonsmooth mechanical systems. *Applied Mechanics Reviews*, 55(2):107, 2002.
- [15] Arunkumar Byravan and Dieter Fox. Se3-nets: Learning rigid body motion using deep neural networks. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 173–180. IEEE, 2017.
- [16] Sylvain Calinon, Florent Guenter, and Aude Billard. On learning, representing, and generalizing a task in a humanoid robot. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(2):286–298, 2007.
- [17] Anindya Chatterjee and Andy Ruina. A new algebraic rigid body collision law based on impulse space considerations. *ASME Journal of Applied Mechanics*, 65(4):939–951, Dec 1998.
- [18] Anindya Chatterjee and Andy Ruina. *A New Algebraic Rigid Body Collision Law Based on Impulse Space Considerations*, volume 65. publisher=MIT press Cambridge, Dec 1998.
- [19] Nikhil Chavan Dafle and Alberto Rodriguez. Prehensile Pushing: In-hand Manipulation with Push-Primitives. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6215 – 6222, 2015.
- [20] Peter Dayan and Geoffrey E Hinton. Feudal reinforcement learning. In *Advances in neural information processing systems*, pages 271–278, 1993.
- [21] Alessandro De Luca and Raffaella Mattone. Sensorless robot collision detection and hybrid force/motion control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 999–1004, 2005.
- [22] Thomas G Dietterich. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of Artificial Intelligence Research*, 13:227–303, 2000.
- [23] Predrag D Djurdjevic and Manfred Huber. Deep belief network for modeling hierarchical reinforcement learning policies. In *2013 IEEE International Conference on Systems, Man, and Cybernetics*, pages 2485–2491. IEEE, 2013.
- [24] Evan Drumwright and Dylan A Shell. A robust and tractable contact model for dynamic robotic simulation. In *ACM symposium on Applied Computing*, pages 1176–1180. ACM, 2009.

- [25] Evan Drumwright and Dylan A. Shell. Modeling contact friction and joint friction in dynamic robotic simulation using the principle of maximum dissipation. In *Proc. of Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2010.
- [26] M. Erdmann. Observing pose and motion through contact. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 723–729, 1998.
- [27] N Fazeli, M Oller, J Wu, Z Wu, JB Tenenbaum, and A Rodriguez. See, feel, act: Hierarchical learning for complex manipulation skills with multisensory fusion. *Science Robotics*, 4(26):eaav3123, 2019.
- [28] Nima Fazeli, Elliott Donlon, Evan Drumwright, and Alberto Rodriguez. Empirical evaluation of common contact models for planar impact. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3418–3425, May 2017.
- [29] Nima Fazeli, Roman Kolbert, Russ Tedrake, and Alberto Rodriguez. Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact. *The International Journal of Robotics Research*, 2016.
- [30] Nima Fazeli, Russ Tedrake, and Alberto Rodriguez. Identifiability analysis of planar rigid-body frictional contact. *The International Symposium of Robotics Research (ISRR)*, 2016.
- [31] Nima Fazeli, Samuel Zepolsky, Evan Drumwright, and Alberto Rodriguez. Fundamental limitations in performance and interpretability of common planar rigid-body contact models. *CoRR*, abs/1710.04979, 2017.
- [32] Nima Fazeli, Samuel Zepolsky, Evan Drumwright, and Alberto Rodriguez. Fundamental limitations in performance and interpretability of common planar rigid-body contact models. In *ISRR*, 2017.
- [33] Nima Fazeli, Samuel Zepolsky, Evan Drumwright, and Alberto Rodriguez. Learning data-efficient rigid-body contact models: Case study of planar impact. *arXiv preprint arXiv:1710.05947*, 2017.
- [34] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [35] M. Gautier and W. Khalil. On the identification of the inertial parameters of robots. In *IEEE Conference on Decision and Control*, pages 2264–2269, 1988.
- [36] Andrew Gelman, Hal S Stern, John B Carlin, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 2013.
- [37] Mevlana C Gemici and Ashutosh Saxena. Learning haptic representation for manipulating deformable food objects. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 638–645. IEEE, 2014.

- [38] Suresh Goyal. *Planar Sliding of a Rigid Body with Dry Friction: Limit Surfaces and Dynamics of Motion*. PhD thesis, Cornell University, 1988.
- [39] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. Planar Sliding with Dry Friction Part 1 . Limit Surface and Moment Function. *Wear*, 143:307–330, 1991.
- [40] Thomas L Griffiths and Mark Steyvers. Finding scientific topics. *Proceedings of the National academy of Sciences*, 101(suppl 1):5228–5235, 2004.
- [41] Heinrich Hertz. On the contact of elastic solids. *J. Reine Angew Math.*, 92:156–171, 1881.
- [42] Francois Hogan and Alberto Rodriguez. Feedback Control of the Pusher-Slider System: A Story of Hybrid and Underactuated Contact Dynamics. In *Workshop on Algorithmic Foundation of Robotics (WAFR)*, 2016.
- [43] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Learning composable models of parameterized skills. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 886–893. IEEE, 2017.
- [44] T. R. Kane and D. A. Levinson. *Dynamics: Theory and Applications*. McGraw-Hill, New York, 1985.
- [45] Frank C Keil. Cognitive science and the origins of thought and knowledge. *Handbook of child psychology: Vol. 1. Theoretical models of human development*, pages 341–413, 1998.
- [46] Charles Kemp, Amy Perfors, and Joshua B Tenenbaum. Learning overhypotheses with hierarchical bayesian models. *Developmental science*, 10(3):307–321, 2007.
- [47] Pradeep Khosla and Takeo Kanade. Parameter identification of robot dynamics. In *IEEE Conference on Decision and Control*, pages 1754–1760, 1985.
- [48] Alina Kloss, Stefan Schaal, and Jeannette Bohg. Combining learned and analytical models for predicting action effects. *arXiv preprint arXiv:1710.04102*, 2017.
- [49] Alina Kloss, Stefan Schaal, and Jeannette Bohg. Combining learned and analytical models for predicting action effects. *CoRR*, abs/1710.04102, 2017.
- [50] Roman Kolbert, Nikhil Chavan Dafe, and Alberto Rodriguez. Experimental validation of contact dynamics for in-hand manipulation. In *International Symposium on Experimental Robotics (ISER)*, 2016.
- [51] Svetoslav Kolev and Emanuel Todorov. Physically consistent state estimation and system identification for contacts. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)*, pages 1036–1043, 2015.

- [52] Michael Koval, Nancy Pollard, and Siddhartha Srinivasa. Pose estimation for planar contact manipulation with manifold particle filters. *The International Journal of Robotics Research*, 7(34):922–945, 2015.
- [53] P. R. Kraus and V. Kumar. Compliant contact models for rigid body collisions. In *Proceedings of International Conference on Robotics and Automation*, volume 2, pages 1382–1387 vol.2, Apr 1997.
- [54] Scott Kuindersma, Robin Deits, Maurice Fallon, Andrés Valenzuela, Hongkai Dai, Frank Permenter, Twan Koolen, Pat Marion, and Russ Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots*, 40(3):429–455, 2016.
- [55] HM Lankarani and PE Nikravesh. A contact force model with hysteresis damping for impact analysis of multibody systems. *Journal of mechanical Design*, 112(3):369–376, 1990.
- [56] Zhe Li and Sridhar Kota. Virtual prototyping and motion simulation with adams. *J. Comput. Inf. Sci. Eng.*, 1(3):276–279, 2001.
- [57] K. M. Lynch and M. T. Mason. Stable Pushing: Mechanics, Controllability, and Planning. *The International Journal of Robotics Research*, 15(6):533–556, December 1996.
- [58] Lucas Manuelli and Russ Tedrake. Localizing External Contact Using Proprioceptive Sensors: The Contact Particle Filter. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [59] D. W. Marhefka and D. E. Orin. A compliant contact model with nonlinear damping for simulation of robotic systems. *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*, 29(6):566–572, Nov 1999.
- [60] Filipe Marques, Paulo Flores, JC Pimenta Claro, and Hamid M Lankarani. A survey and comparison of several friction force models for dynamic analysis of multibody mechanical systems. *Nonlinear Dynamics*, 86(3):1407–1443, 2016.
- [61] Matthew T Mason. Toward robotic manipulation. *Annual Review of Control, Robotics, and Autonomous Systems*, 1:1–28, 2018.
- [62] Brian Mirtich. *Impulse-based Dynamic Simulation of Rigid Body Systems*. PhD thesis, University of California, Berkeley, 1996.
- [63] Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 3303–3313, 2018.
- [64] Victor Ng-Thow-Hing and Petros Faloutsos. Dance: Dynamic animation and control environment. 1999.

- [65] Albert Nubiola and Ilian A Bonev. Absolute calibration of an abb irb 1600 robot using a laser tracker. *Robotics and Computer-Integrated Manufacturing*, 29(1):236–245, 2013.
- [66] Jong-Shi Pang and Jeffrey C. Trinkle. Complementarity Formulations and Existence of Solutions of Dynamic multi-Rigid-Body Contact Problems with Coulomb Friction. *Mathematical Programming*, 73(2):199–226, 1996.
- [67] Ronald Parr and Stuart J Russell. Reinforcement learning with hierarchies of machines. In *Advances in neural information processing systems*, pages 1043–1049, 1998.
- [68] Jim Pitman. Combinatorial stochastic processes. *Technical Report No. 621*, 0(August), 2002.
- [69] R Platt and L Kaelbling. Efficient planning in non-gaussian belief spaces and its application to robot grasping. *International Symposium on Robotics Research (ISRR)*, 2011.
- [70] Robert Platt Jr, Russ Tedrake, Leslie Kaelbling, and Tomas Lozano-Perez. Belief space planning assuming maximum likelihood observations. 2010.
- [71] M. Posa, C. Cantu, and R. Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 33(1):69–81, October 2013.
- [72] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.
- [73] Alberto Rodriguez and Matt T. Mason. Failure detection in assembly: Force signature analysis. In *Automation Science and Engineering (CASE), 2010 IE*, pages 210–215. IEEE, 2010.
- [74] O.S. Salawu. Detection of structural damage through changes in frequency:A review. *Engineering Structures*, 19(9):718–723, 1997.
- [75] Tanner Schmidt, Richard A Newcombe, and Dieter Fox. Dart: Dense articulated real-time tracking. 2014.
- [76] J.-J. E. Slotine. On the Adaptive Control of Robot Manipulators. *The International Journal of Robotics Research*, 6(3):49–59, 1987.
- [77] Peng Song, Peter Kraus, Vijay Kumar, and Pierre Dupont. Analysis of rigid body dynamic models for simulation of systems with frictional contacts. *J. Applied Mech.*, 68, 2000.
- [78] David E. Stewart. Rigid-Body Dynamics with Friction and Impact. *SIAM Review*, 42(1):3–39, January 2000.

- [79] David E. Stewart and Jeff C. Trinkle. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *Intl. J. Numerical Methods in Engineering*, 39(15):2673–2691, 1996.
- [80] D.E. Stewart and Jeffrey C. Trinkle. An Implicit Time-Stepping Scheme for Rigid Body Dynamics with Inelastic Collisions and Coulomb Friction. *International Journal for Numerical Methods in Engineering*, 39(15):2673–2691, 1996.
- [81] W. J. Stronge. Rigid body collisions with friction. *Proc. of the Royal Society of London A*, 431(169–181), 1990.
- [82] Richard S Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- [83] David Tax, Alexander Ypma, and Robert Duin. Support Vector Data Description Applied to Machine Vibration Analysis. In *5th Annual Conference of the Advanced School for Computing and Imaging*, pages 398–405, 1999.
- [84] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 5026–5033. IEEE, 2012.
- [85] Marc Toussaint, Kelsey Allen, Kevin A Smith, and Joshua B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning.
- [86] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. *arXiv preprint arXiv:1809.10790*, 2018.
- [87] Y. Wang and M. T. Mason. Two-dimensional rigid-body collisions with friction. *ASME J. Appl. Mech.*, 59:635–642, 1992.
- [88] Zi Wang, Caelan Reed Garrett, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Active model learning and diverse action sampling for task and motion planning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4107–4114. IEEE, 2018.
- [89] E. T. Whittaker. *A Treatise on the Analytical Dynamics of Particles and Rigid Bodies, 4th ed.* Dover, 1944.
- [90] A.S. Willsky. A Survey of Design Methods for Failure Detection in Dynamic Systems. *Automatica*, 12(6):601–611, 1976.
- [91] Kuan-Ting Yu, Maria Bauza, Nima Fazeli, and Alberto Rodriguez. More than a Million Ways to be Pushed. A High-Fidelity Experimental Data Set of Planar Pushing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.

- [92] Kuan-Ting Yu, John Leonard, and Alberto Rodriguez. Shape and Pose Recovery from Planar Pushing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1208 – 1215, 2015.
- [93] Li Zhang, Siwei Lyu, and Jeff Trinkle. A dynamic bayesian approach to real-time estimation and filtering in grasp acquisition. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 85–92, 2013.
- [94] Li Zhang and Jeffrey C. Trinkle. The application of particle filtering to grasping acquisition with visual occlusion and tactile sensing. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3805–3812, 2012.
- [95] Jiaji Zhou, Robert Paolini, J. Andrew Bagnell, and Matthew T. Mason. A convex polynomial force-motion model for planar sliding: Identification and application. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [96] Jiaji Zhou, Robert Paolini, Aaron M Johnson, J Andrew Bagnell, and Matthew T Mason. A probabilistic planning framework for planar grasping under uncertainty. *IEEE Robotics and Automation Letters*, 2(4):2111–2118, 2017.