Modeling, Planning, and Control for Whole-Body Manipulation of Unknown Objects with Large-Scale Soft Robots

by

Curtis C. Johnson

A prospectus submitted to the faculty of
Department of Mechanical Engineering
Brigham Young University

October 24, 2022

Prospectus Approval

Prospectus submitted by:

Curtis C. Johnson | Date

This prospectus has been approved by each member of the Graduate Committee:

Committee Member - Chair | Date

Committee Member | Date

Committee Member | Date

Committee Member | Date

1 Problem Statement

The term 'manipulation' refers to an agent's control of its environment through selective contact [1]. Humans and animals are expert open-world manipulators. The term 'open-world' refers to the idea that the world is inherently unstructured and objects inside of it have infinite variability [2]. Humans are capable of manipulating a wide range of objects varying in size, shape, texture, stiffness, and mass in unstructured environments – often without knowing anything about the object beforehand. While many manipulation tasks can be completed using only hands, many tasks require the use of additional body parts, such as the arms, chest, legs, back, or shoulders. These situations arise when the object is particularly large, heavy, or awkward. For example, a delivery person can move boxes of all shapes and sizes from a truck to a house; a firefighter can lift and carry someone to safety; a nurse can help someone out of bed. These are also cases where a robotic assistant could be beneficial for the safety and well-being of the human, whatever their role in the task.

Unfortunately, many state-of-the-art robotic manipulators are practically useless in these situations, as they are tailored towards the exclusive use of end effectors to accomplish a task. Entire books have been written with this end-effector-only perspective. But what if a robot is tasked with carrying a disaster victim to safety or delivering a large box? What should the robot grasp with its end-effector? In some instances customized hardware solutions are available (like suction grippers used in the logistics industry to lift heavy boxes). While effective, developing customized hardware solutions for every use case is prohibitively expensive for a general-purpose assistive robot.

Tasks such as these require a less restrictive approach to manipulation. What if the robot could use any part of its structure or the environment to perform a manipulation task, instead of using only well-defined contact locations? Suddenly a considerable quantity of contact-rich manipulation behaviors like enveloping, bracing, squeezing, sliding, or supporting can be exploited in control and planning algorithms. Manipulation of this kind is referred to as whole-arm, whole-body, or whole-world manipulation.

The concept of whole-arm manipulation was originally proposed in the 1980s by Salisbury [3] with the WAM arm along with basic kinematic and dynamic requirements for making and maintaining contact with manipulated objects. Since then, the majority of the research in whole-arm manipulation has relied on carefully tuned, task-specific impedance control or force control, neither of which offer robustness for handling open-world scenarios. Often accompanying these controllers is an underlying assumption of one or two contact points in specific locations. Additionally, the manipulation is generally assumed to be quasi-static and objects are assumed to have known inertial properties and states that are both certain and fully-observable. These assumptions are unrealistic for open-world manipulation.

Meanwhile, recent innovations in the field of soft robotics suggest that purposefully incorporating passive compliance into the manipulator can dramatically simplify the control problem while operating around people or in unstructured environments. Passive compliance makes a robot inherently safe and robust in these open-world situations. Rigid robots require sophisticated control and expensive sensors to accomplish the same level of robustness and safety. Using passive compliance can mitigate the risk of contact rich tasks and can also help simplify the problem.

For my thesis I propose to develop control and planning algorithms for autonomous whole-

Figure 1: Rendering of preliminary robot torso design. This will be used for hardware testing of whole-arm manipulation tasks in this thesis. For scale, the entire structure is about 1.8 meters tall.

body manipulation of objects that cannot be grasped with an end effector. The proposed algorithms will use only onboard sensory feedback (i.e. tactile, pressure, orientation, and visual feedback) to 1) learn inertial and geometric properties of unknown objects and 2) manipulate the object. This sensory constraint reflects the fact that instrumenting every object to be manipulated with various sensors is impractical in open-world manipulation. The algorithms will be tested in a simulation environment, as well as on hardware using a two-armed soft robot torso outfitted with tactile sensors, currently in development (see Figure 1). This hardware platform enables us to also explore the use of passive mechanical compliance and its role in whole-body manipulation.

2 Background

2.1 Challenges of Whole-Body Manipulation

Several textbooks have been written on robotic manipulation ([4], [5], [6], [7], [8]), but in general these resources, along with current research, are narrowly focused on gripper-only manipulation. My work will focus specifically on whole-body manipulation, a type of manipulation that is not well-studied in the literature–especially in regards to soft robots–but which has a huge potential value to the field. This section is organized according to the main challenges presented in [1], that make robotic manipulation difficult: mechanisms, perception, modeling and control with contact, and planning.

2.1.1 Mechanisms

The challenge of mechanism design lies inherently in the tension between multiple desirable attributes of a general-purpose manipulator: safe in unstructured environments, reasonably strong and fast, and relatively inexpensive. Successful manipulation requires that all three of these requirements be reasonably met.

There are two different directions of research attempting to satisfy all three requirements. The first is outfitting normal rigid robots with soft materials. The authors of [3] first recommended covering the rigid structure with high-friction, soft materials. The authors of [9] use a similar approach with pressurized foam pads on the chest and arms of a rigid robot torso for whole-arm bimanual grasping. Many researchers simply add soft pads at the expected contact locations, for example in the bimanual manipulation of a box [10] or assistance of a hospital patient [11].

Another approach is to soften the body of the manipulator itself. This can be done by adding compliance into the joints of an otherwise rigid robot, as is done with series elastic actuators. The field of soft/continuum robotics goes even further and reduces or removes typical rigid linkages throughout the entire robot. Many such designs inspired by octopus tentacles, muscles, or elephant trunks have been proposed for whole-arm manipulation [12], [13]. There is also a related body of research in the field of robotic hand design [14], [15] which borrows some of the same ideas for manipulating small objects (e.g. gecko-like grippers [16]).

It is clear that the research community is converging on the idea of mechanical intelligence [17] where mechanism design can help simplify the contact-rich manipulation problem instead of relying on overly-cautious, avoid-contact-at-all-costs planning and control algorithms.

2.1.2 Perception

The role of perception is often minimized in robotic manipulation research by assuming perfect knowledge of the object and the environment. This assumption necessitates the use of high-accuracy systems such as motion capture and robot-facing cameras to avoid occlusions. But these sensors are not practical for use in the real world, and the control algorithms developed with them are also overly reliant on that type or quality of sensory information. In order for a robotic system to be successful in open-world manipulation of an unknown object, the perception subsystem must be physically plausible and needs to provide both global and local information about the manipulation task at a sufficiently high rate and resolution to be useful. But the local information cannot be isolated to a few known points, since contact can occur anywhere along the structure of the robot during whole-arm manipulation. The authors of [18] provide a useful framework with which we can categorize these pieces of information, namely: contact-level, object-level, and action-level information. Contact-level information encompasses any local information at a particular contact site (e.g. texture, or forces). Object-level information includes information inherent to the entire object and is not restricted to a certain location (e.g. shape or inertial properties). Action-level information is used for control and planning sequences of actions (e.g. object pose). Contact-level and action-level information are summarized in this section. Section 2.2 will provide an in-depth discussion of object-level information.

Onboard vision systems can provide reliable action-level information, and are often used very successfully in manipulation in the context of reinforcement learning ([19], [20]). But these approaches suffer from problems like occlusions, and also cannot provide adequate contact-level information that is important for large/heavy objects.

There are many localized tactile sensing solutions which provide contact-level information for in-hand manipulation [21]. Force-torque sensors provide excellent contact-level information ([22], [10]) but they are expensive and limited to single known contact points, making them unsuitable for whole-arm manipulation tasks. There are vision and tactile combinations (i.e. visuotactile sensors) [23], but again, these are small-scale vision-based tactile sensors. It is not clear how to scale these technologies for whole-arm manipulation.

Incorporating vision and distributed tactile sensing individually (as opposed to the visuotactile sensors discussed above, where both methods are embedded into a single device) could allow scaling up to the required physical area needed for whole-body manipulation. Distributed tactile sensing is an active area of research with innovative designs in development ([18], [24], [25]), but the sensory feedback from these designs has not been used meaningfully in control or autonomous manipulation. Additionally, methods to effectively fuse distributed tactile signals with vision have not been investigated (i.e. sensor fusion).

2.1.3 Modeling and Control with Contact

Dynamic simulation in robotics is an extremely useful tool. It provides a means to generate massive amounts of training data for machine learning, allows the use of more sophisticated control algorithms, and accelerates the testing and verification of new algorithms in a safe manner [26]. Desirable features in a simulation environment include simulation speed, parallelizability, physical accuracy, differentiability (for optimization), and realistic rendering (for pixel-space learning), to name a few. There are many high-quality simulators for robotics available, which offer various levels of these features ([27], [28]). One of the most difficult challenges which affect all simulators is the so called 'sim2real' gap, where simulations fail to capture real-world physical phenomena or introduce non-physical artifacts [29]. The sim2real gap is perhaps largest in the context of contact mechanics, which are crucial in whole-arm manipulation. A wide variety of contact models are used in robot simulation ([30], [31]) and it is still unclear which models most accurately capture real-world contact dynamics, despite several published comparison studies ([32], [33]).

Early work on control for whole-arm manipulation [34] explores the idea of line stiffness control which enables the entire link of the robot to search for objects through collision and push them across a table. The authors of [35] and [36] established the basic kinematic and force-based control requirements to make and maintain contact.

Later, the author of [37] developed the 'Octograsp' algorithm for generating whole-arm grasps using motors and joint torque sensors. The algorithm essentially moves the proximal links inwards until contact is detected and continues down the arm until the object is fully enveloped, while maintaining sufficient contact forces. This approach is similar to the approach used in [9] though the method of detecting contact was through soft pressurized sleeves instead of joint torque measurements. This enveloping algorithm simplifies the whole-arm manipulation problem such that an explicit planner is not needed. While authors demonstrate that this method generalizes to several different types of objects, it only allows for planar grasps. It also relies on the assumption that the object is placed in the workspace of the arm and in a pose which allows 'hugging'. The algorithm does not treat re-grasping either, which may be necessary when working in three dimensions, or when learning about an object.

Instead, some force-based control methods explicitly reason about contact. The authors of [38] used model predictive control (MPC) and impedance control with multiple contact locations and were able to successfully regulate contact forces while moving through clutter. While no prior model of the environment was assumed, it was assumed that the task was quasi-static and planar. Later in [39], they extended MPC to a dynamic task with static obstacles. This work illustrates the value of MPC and its ability to deal with constraints explicitly. The authors of [40] instead address the manipulation problem with an object-level impedance controller. They use the concept of virtual springs between three known contact locations to perform a stable whole-arm grasp of a large ball. However, in whole-arm manipulation assuming known contact locations beforehand is unrealistic. While these locations could be measured using some sort of tactile sensing, it is not clear how this method could scale to more uncertain contact locations. Additionally, the physical properties of the object are also assumed to be known since the controller uses feedforward terms to compensate for its dynamics. This makes the control method particularly sensitive to modeling errors.

The authors of [41] present a nonlinear model predictive control (NMPC) method which is solved using a generalized form of an iterative linear quadratic regulator (iLQR). They successfully deploy the controller on quadruped robots trotting and jumping, and report control rates of up to 190 Hz. The authors of [42] propose a formulation of MPC which uses a mixed-integer quadratic program (MIQP) to handle contacts of a cart-pole system colliding with soft walls as discrete optimization constraints. The authors of [10] use a slightly different approach by using MPC in contact space (as opposed to joint space) to control the deformation trajectory of the soft pads that they mounted to the robot end effector. All of these methods demonstrate various ways to successfully incorporate a few known contacts into an MPC-based scheme.

In summary, researchers often make one of two large assumptions to solve the manipulation problem. One assumption is that knowing explicit information about the object is unnecessary; the other is that information (e.g. geometric, inertia, etc.) about the manipulated object is entirely known. The first assumption limits manipulation capabilities to semi-static grasping, while the second clearly cannot be the case for open-world manipulation as it would require building a library of thousands of objects, all with different properties. Additionally, it is unclear from these approaches how to formulate and/or scale the optimal control problem when contact can be made at multiple locations.

2.1.4 Planning

Three major approaches to planning for whole-arm manipulation in the current literature are learning, graph/tree search, and offline trajectory optimization in simulation.

The authors of [43] approach the problem as a continuous-state Markov Decision Process (MDP) and use reinforcement learning to approximate a Q-value function and obtain an optimal policy. The authors report learning and planning a successful planar whole-arm manipulation policy in under two seconds. The major challenges reported by the authors are that the MDP framework used 1) does not scale well to higher dimensions, which will occur for three-dimensional manipulation; 2) is sensitive to uncertainty of the world state; and 3) requires fairly accurate measurements of every object.

To address the need for accurate object measurements beforehand, the authors of [44] instead explore learning how to grasp without seeing. In this work, they use a touch-based object localization algorithm to generate an initial grasp which is improved upon as an autoencoder neural network learns to recognize the material. Though this gripper-only, planar approach does not scale to three-dimensional whole-arm manipulation, the work demonstrates the potential benefits of tactile exploration in whole-arm manipulation planning.

Authors of [45] present some dimensionality reduction techniques for model-free reinforcement learning of bi-manual, three-dimensional, (gripper-only) manipulation in unstructured environments. The techniques include the use of various types of joint-space movement primitives to efficiently encode a trajectory with a limited set of parameters [46]. These movement primitives can then be combined and adapted to obtain more complex motions. A similar idea called 'manipulation primitives' exists for efficiently parameterizing trajectories in contact space, as is done with tactile feedback in [47] and for robotic assembly [48]. The use of primitives has not been explored in context of three-dimensional whole-arm manipulation, with only a few planar proof-of-concept implementations, as in [9].

The use of manipulation primitives lends itself well to representing the manipulation problem as a graph/tree search. The authors of [49] use rapidly exploring random trees (RRTs) to kinematically plan whole-arm grasps of objects tumbling in space. This is done in three dimensions by grasping in a 'twining plane', i.e. a curve that passes through the center of mass of an object and is perpendicular to the object's rotation axis. By planning whole-arm grasps which lie in this plane, the arms can catch the object robustly. This method relies heavily on known inertial parameters before manipulation occurs, but the approach is one of the only works that consider three-dimensional whole-arm manipulation (as opposed to strictly planar manipulation). The authors of [50] also plan three-dimensional whole-body manipulation, but instead use graph searches of manipulation graphs. The manipulation graph is built with nodes representing object poses, and transitions between nodes representing manipulation primitives such as pushing, rotating, lifting, etc. When the robot detects that the planned object motion does not occur, it searches for a new plan along a different manipulation graph whose parameters match the real world more accurately. This method successfully generates complex manipulation strategies that are robust and adaptive, but relies heavily on having several manipulation graphs for each object and therefore does not scale well to general manipulation.

The last of the three general approaches is offline trajectory optimization. These methods consider the problem of optimizing state, input, and contact trajectories simultaneously over some future horizon, an idea now called 'contact-implicit' trajectory optimization. The general approach is to use some simulation model to optimize a trajectory and then follow the trajectory with PID control loops in hardware. There are various ways to actually perform the optimization depending on the formulation of the problem. The authors of [51] and [52] both use an unconstrained iLQR algorithm. The authors of [53] formulate a nonlinear optimization with complementarity constraints and use sequential quadratic programming (SQP). Several other works have built on this idea to improve accuracy via orthogonal collocation [54], by using ideas from discrete variational mechanics [55], or by using different contact models [52]. These new contact-implicit trajectory optimization methods are promising, especially for multiple contacts, but have largely only been tested in simulation. Very few trajectory optimization techniques have been extended to real-time control for robustness to uncertainty (e.g. by using model predictive control) or to whole-arm manipulation.

2.2 Challenges of Object Learning

Object learning is a central component of autonomous robotic manipulation. My work will focus specifically on the manipulation of 'unknown objects', defined as objects which have both geometric and physical properties that are uncertain [56]. As was briefly discussed in Section 2.1.2, vision is one of the most ubiquitous on-board sensors that can be used for manipulation, and has been used successfully. But cameras alone cannot provide information on inertial properties–which are important for whole-arm manipulation–and cameras suffer from other problems like occlusion, which will happen frequently during this type of manipulation. As such, this thesis will explore the complementary properties of tactile sensing and vision for whole-arm manipulation.

Depending on the type of manipulation, the robot may need only rough geometry estimates to successfully manipulate the object (such as is often the case with finger-tip grasps, where the weight of object is negligible compared to grasp forces). This assumption cannot be made for whole-arm manipulation. Whole-arm manipulation is most useful for objects that are heavy, awkward, or large. In these cases, it is important to know both geometric and inertial properties of the object. Many aforementioned papers on control/planning for robotic manipulation simply assume that these properties are known beforehand. In open world manipulation, this is not be the case. The algorithms developed in this thesis will not make that assumption, and will therefore rely on simple onboard sensors to learn the necessary information for whole-arm manipulation. This section will folow the outline presented in [18] which categorizes the different types of object-level information that are commonly needed for robotic manipulation into three main categories: object localization, object shape, and object mass and dynamics. In theory, a spectrum of object learning methods could exist between vision only and tactile feedback only. Yet the authors of [57] outline many object learning methods that are either vision-based or exploratory (i.e. via tactile feedback). Curiously, there is no mention of combining these two methods, leaving a clear area for potential contributions.

2.2.1 Object Localization and Shape Estimation

Object localization and object shape estimation is almost exclusively done with vision via RGBD sensors. Various algorithms exist for processing RGBD sensor data streams to accomplish object detection, tracking, and shape estimation (see Chapters 4-6 of [2]). Unfortunately, these algorithms suffer heavily from issues related to occlusion. To illustrate, when using these algorithms it is not uncommon to find 8-10 cameras–both on and off the robot–focused on a specific environment to avoid occlusion problems. This is clearly not scalable to open-world manipulation.

The growth of deep learning has had a large impact on robotic perception, in this case with the use of convolutional neural networks (CNNs). These networks use RGBD sensory inputs to provide important imformation on object pose and shape [56]. For example, [58] uses a CNN originally presented in [59] called ShapeHD to perform 'shape completion'. Vision feedback from a single camera provides a two dimensional representation of the object which is often occluded. The ShapeHD CNN maps this to a three-dimensional point cloud, with which the robot can reason about possible grasps. These methods help to solve the anticipated occlusion problem for whole-arm manipulation, but still do not provide any inertial parameters. In fact, because many of these algorithms were developed for in-hand manipulation of relatively small and light objects, they often do not consider physical properties of the object at all.

There is another branch of research working toward combining vision and tactile based sensing for shape estimation [60], [61]. These sensors use a deformable gripper material with an embedded camera to estimate deformation patterns as a grasp is executed. These sensors can provide all of the feedback types that are required for whole-arm manipulation (estimate object pose, geometry, and mass properties), but are limited to fingertip grasps as the embedded camera's field of view constrains the possible size of the deformable material.

2.2.2 Physical Parameters Estimation

While many manual methods exist for accurately estimating the inertial parameters of rigid bodies (e.g. scales, bifilar pendulums, etc.), they have not been extended for autonomous robotic manipulation. The simplest extension of these methods is to measure the resultant wrench vector using a force/torque sensor on the end-effector of the robot. As previously discussed in Section 2.1.2, this cannot scale to whole-arm manipulation.

There are vision-based methods to estimate inertial parameters, but they rely heavily on assumptions like uniform density which are generally not known beforehand. Some recent work [57] leverages video clips of objects in simulation or real life to learn inertial properties of the object in the video.

While it does require more hardware than vision, physical exploration of an object with tactile sensing allows much more accurate estimation of inertial parameters without needing large datasets and strong prior assumptions. The authors of [57] outline many tactile exploration methods such as tilting [62], pushing, lifting, or swinging [63] to estimate inertial parameters of both small and large objects. A major disadvantage of these approaches is that they require a time consuming process of optimizing touch sequences to get the required data [18].

One way around this is presented in [64], where estimation and motion planning happen simultaneously for a humanoid robot manipulating a large box. Their algorithm recursively estimates inertial parameters with Bayesian methods while planning the feasibility of manipulation primitives based on current likelihood of physical parameters. Thus, the robot can start manipulating the object while learning about it, instead of waiting to know everything. But this ability comes at a high computational cost. In response to this, the authors note that the approach is limited in its ability to scale to real-world (and more specifically, real-time) robotic manipulation.

2.3 Challenges of Manipulation with Large-Scale Soft Robots

The first main challenge preventing the use of soft robots in manipulation is the difficulty of scaling up soft robots. State-of-the-art soft robot platforms capable of manipulation are typically made from various types of silicone which are are quite dense–the density of silicone rubber (1000-1500 kg/m^3) is more than the density of water. Robots made of these materials are inherently limited in size, as larger designs would likely be unable to lift their own weight. Some recent design innovations have lead to larger, meter-scale soft robots ([65], [66]), which are capable of lifting large objects and walking. Very little has been done in regards to soft robots manipulating objects, likely because the payload of most soft robots scale poorly. The most promising designs for manipulation are the so-called continuum manipulators inspired by nature's masterful manipulators like the octopus or the elephant. These types of robots constitute a middle ground between silicone-based soft robots and rigid robots. They are comprised of soft segments interconnected with rigid components and are driven pneumatically or with cables and motors. As such, these types of robots are typically capable of higher payloads while still keeping the desirable trait of natural compliance, and are good candidates for whole-arm manipulation of large objects. These useful properties have already been exploited to estimate object shape by wrapping [67] and for planning while in contact with the environments [68].

The second main challenge is the lack of accurate models and high-performance controllers for soft robots. There are many different modeling approaches ranging in computation time and complexity (Cosserat rods, FEM, rigid body approximations, machine learning, etc), but the the underlying uncertainties of soft robots present a problem that has still not been solved. Uncertainties can come from materials, manufacturing, wear over time, kinematics, etc. Some have taken a data driven approach to either learn a model [69] or learn an effective control law directly (e.g. via reinforcement learning [70]). These approaches have demonstrated good results. But the time required to gather sufficient data is still large and it is unclear what types of data lead to high-performing models. In any case, there is also the problem that the process of gathering this data can wear out a soft robot. Exacerbating the data collection problem, particularly in a manipulation context, is the fact that because of the deformation inherent in soft robots, different loading conditions will cause the dynamics to change. Gathering a sufficient amount of representative data to accurately model these effects across all possible loading conditions is very challenging and expensive.

Instead of gathering data directly on hardware, using simulation tools to generate training data is a promising alternative. However, the aptly named 'sim2real gap' (i.e. the differences between a simulation and the real world) is still quite large for soft robots. Generating physically accurate data in a computationally feasible way is an active area of research with some promising results, but it is still not clear how to best close the gap.

Despite their challenges there are clear benefits of using simulation for data collection, planning, and control. Because of the increased availability of affordable computation power, simulations are becoming increasingly important in robotics. But common tasks required of robotic manipulators involve complicated frictional contacts which are nonlinear, non-smooth, and difficult to simulate quickly and accurately. Also important to note is that in unstructured environments where soft robots are likely to be most effective, the number, position, and magnitude of such contacts is also generally unknown (or at least uncertain).

Simulating soft robots in such environments generally has two main considerations: the simulation of the robot itself and the simulation of contacts.

Until recently, the simulation of soft robots has generally been comprised of project-specific simulation environments, but that limits speed of development and makes it harder to benchmark algorithms. The problem is that the options commonly available for simulation of soft robots are either FEM or rigid-body simulations. Some recent work has attempted to solve this problem. The authors of [71] developed SoMo, which is a simulation framework specifically for soft continuum robots built on top of Bullet, a rigid-body physics engine. The authors of [72] use Drake [73], a rigid-body dynamics library, with an augmented rigid body model [74] to simulate a soft continuum manipulator, but they could not simulate the model in real time because a solution to a large set of nonlinear equations was required. ChainQueen ([75],[76]) and SOFA [77] are two simulators built specifically for soft robots, but ChainQueen does not support rigid-soft hybrid robots, which we hypothesize are important for contact tasks and payload capabilities, and SOFA relies on FEM which is currently too slow for real-time control.

3 Research Objectives

The goal of this research is to enable autonomous whole-body manipulation of large objects. Success means that a soft robot will be able to learn about an unknown object in order to move it into a desired configuration. To this end, the main objectives of the proposed research are as follows:

- Develop dynamic models and evaluate a simulation framework of the proposed soft robot torso with a particular focus on dynamics and contact with the environment. This will be used for testing algorithms throughout the work.
- Explore the combined use of visual and tactile sensory feedback and develop algorithms which use that feedback to learn geometric and inertial properties of large unknown objects.
- Develop and evaluate algorithms that utilize learned knowledge about a large object in order to manipulate it into a desired configuration.

These objectives will advance the state-of-the-art by significantly extending the manipulation capabilities of robots as well as providing algorithms more directly applicable to the real world, where objects cannot all be known and localized beforehand. These algorithms, though specifically developed to take advantage of the unique properties of soft continuum manipulators, will provide valuable insight for the robotics community in general. The workspace of the the robot is effectively extended if it can use its whole arm, the payload is increased because loads can be supported by the structure of the robot, and the amount of objects that a robot can manipulate is much larger.

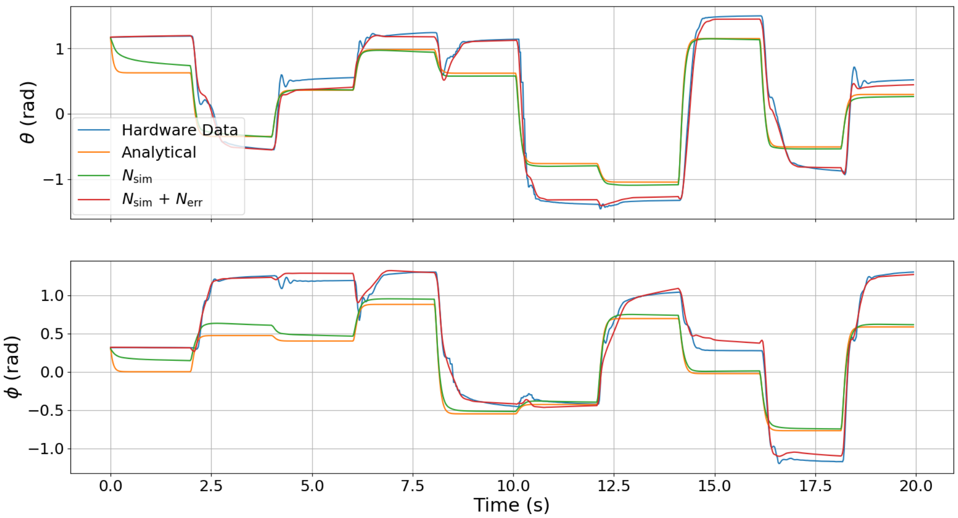

Figure 2: Joint angle predictions from [69]. The hardware data is the raw data collected on the robot. The analytical model derived in [78] clearly is not accurate. The $N_{\text{sim}} + N_{\text{err}}$ line is the model after learning took place.

4 Proposed Research

The following sections expand on the specific technical approach I plan to follow to accomplish the research objectives in Section 3. Note that some of the work I will describe has already been published. I will be using a soft robot torso (see Figure 1) for hardware experiments throughout this work. Each arm will be outfitted with fabric tactile sensor arrays adapted from [25].

4.1 Dynamic Modeling and Simulation of Soft Robots

As has been discussed, soft continuum robots have some inherent advantages (i.e. 'mechanical intelligence') that can make contact-rich manipulation an easier problem. These advantages do come at a cost; soft robots are difficult to model accurately. In addition, state-of-the-art simulators for robotics generally do not support soft robots.

In my previously published work [69], we used deep learning to compensate for the modeling errors that we knew existed in the dynamic model originally developed in [78]. The resulting model (with error compensation) more accurately predicted the states of the real robot (see Figure 2) and improved control performance.

While effective, this approach was difficult to scale to more degrees of freedom. In [79], we solved this problem by using the Recursive-Newton-Euler (RNE) algorithm to simulate the dynamics numerically instead of analytically. Each continuum joint is considered a chain of thin disks, each of which has a mass and inertia. The RNE algorithm simply steps along each disk and computes relative velocities, accelerations, and forces. This method worked well, but was another instance of an ad-hoc soft robot simulation. It also would have required custom implementation of contact and sensor models. No off-the-shelf simulator exists that does both soft/rigid body simulation and contact quickly enough or accurately enough.

My proposed approach builds on these ideas, but uses a well-supported rigid-body physics simulator called MuJoCo (Multi-Joint Dynamics with Contact) [80]. MuJoCo uses RNE for dynamic simulation, but also implements a state-of-the-art contact model along with many other types of sensors. By treating each of the disks as a distinct rigid body interconnected by universal joints with stiffness and damping, we can simulate the continuum nature of the joints, as shown in Figure 3 for a single arm. MuJoCo, similar to many other robot simulator frameworks, uses an XML-based modeling language. This means that the models created for this proposed research will be portable, self-contained, and compatible with other research to allow for easier comparison–a major improvement over the ad-hoc simulations currently available for soft robots.

The idea of approximating a flexible body with smaller rigid bodies interconnected with flexible links is not novel. It is very similar to the pseudo rigid-body model originally introduced for compliant mechanisms in 1996 [81], and to what was recently implemented for planar bending of soft robots in [71]. My work will extend the idea to three dimensions while building on top of well-supported, standard tools that are widely used in the robotics community. The proposed simulator will also accelerate development of the algorithms described in the following sections.

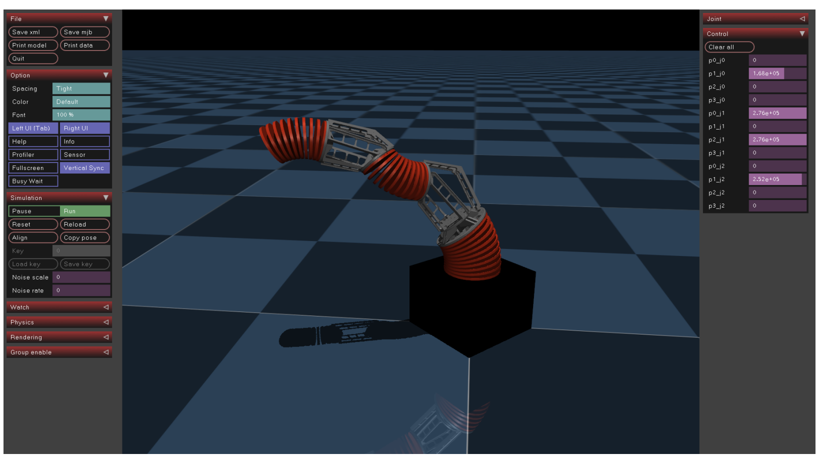

Figure 3: Screenshot of continuum soft robot simulation. This model contains only one arm, but the final simulation will include two arms with straight links and a torso, as in Figure 1.

4.2 Learning about Unknown Objects

My approach to accomplish autonomous learning about unknown objects will focus particularly on combining distributed tactile sensing with vision. Recall that an unknown object is defined in this work as an object with uncertain geometric and physical properties.

The first objective to learning about an unknown object is to obtain an estimate of the object's geometry. I propose to build on [82], which uses a convolutional neural network to perform shape completion (i.e. predict a volumetric mesh of an occluded object) using RGBD sensory input. The authors have open-sourced their model and training data, making this a good starting point. Because their work only considered hand-held objects, it is not clear how

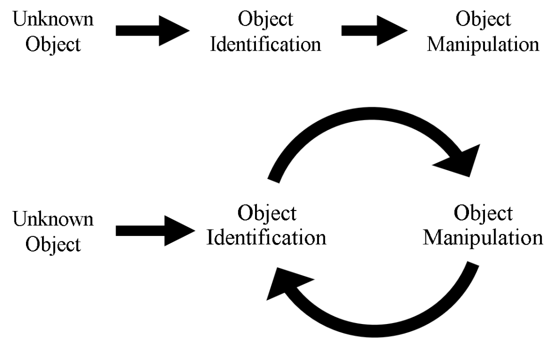

Figure 4: Different approaches to manipulation of an unknown object. The top row is the general assumption that an object must be fully identified before manipulation can take place. The bottom row is a different perspective, where manipulation and identification help improve one another iteratively.

the pre-trained model will generalize to larger objects typical of whole-body manipulation. Part of my work in this area will be to evaluate the performance of the pre-trained model and fine-tune it with custom data for larger objects if necessary.

Once a preliminary geometric model is found, the remaining task is to learn about the physical features important for whole-body manipulation. Because of the lack of literature on whole-arm manipulation, it is not clear which physical parameters will be most useful. The authors of [63] found that surface friction, object mass, center of mass, and the moments of inertia are the most important for an object swing-up task. These same properties are likely useful for whole-arm manipulation, so my work in this area will initially focus on developing algorithms to estimate these values autonomously.

As was discussed in Section 2.2, estimating these parameters with vision sensors alone requires strong assumptions, and estimating these properties with tactile sensors alone requires an expensive process of optimizing and performing a sequence of exploratory actions. The relatively unexplored area of combining these sensors will be a major focus of my work. I propose a method which uses the strong assumptions generally used with vision sensors (i.e. uniform density) as a prior, which will be iteratively refined using a limited sequence of exploratory actions. While assumptions like uniform density are not valid in general, especially for large or awkward objects, it can still serve as a decent 'warm start' solution which can dramatically increase the efficiency of any necessary exploratory actions that follow.

The exploratory actions can then be designed or learned from movement primitives such as tilting, swinging, pushing, or hefting. For example, tilting an object in different directions can provide information about the center of mass. Pushing an object can provide information about coefficients of friction. Hefting, or similar dynamic movements, can provide information about the inertial properties. This step of object learning is also a valuable opportunity to explore and exploit the complementary relationship between object learning and object manipulation. The general assumption in the literature is that the object to be manipulated must be fully identified before manipulation can occur. This may not be the case. In fact, control and planning during manipulation is likely necessary for object learning and vice versa. Figure 4 illustrates the differences between these perspectives.

4.3 Control and Planning for Whole-Arm Manipulation

Despite the best system identification efforts, the model and simulation framework proposed in the previous section will never completely capture reality because of the complexity and uncertainty of soft robots. As a result the controllers and planners used in manipulating an object must be robust and adaptive.

MPC is a promising candidate for whole-arm manipulation because of its ability to explicitly include constraints, its inherent robustness to uncertainty, and its forward prediction capability. In [78], we showed that MPC can be augmented with an adaptation law borrowed from Model Reference Adaptive Control (MRAC). The adaptation law allows the controller to improve in response to tracking error. This was demonstrated on a single 2-DOF soft robot joint and has yet to be scaled to more degrees of freedom. To accomplish robust and adaptive joint-space tracking, I plan to extend this controller to the soft robot torso in Figure 1 which will have at least 13 degrees of freedom. An additional source of improvement can come from the interaction between the object manipulation stage and object identification stage, as shown in Figure 4.

The ability to explicitly include constraints will be valuable when reasoning about contact modes. Because it is likely that contact-level constraints will cause optimization convergence problems, I plan to incorporate the idea of 'contact implicit' trajectory optimization. This was discussed in Section 2.1.4, where the idea was used in offline trajectory optimization for quadrupedal locomotion.

I provide a basic overview of the theory behind contact-implicit optimization here. The forward dynamics of a rigid-body undergoing inelastic contacts can be formulated as a Linear Complementarity Problem (LCP):

$$\begin{aligned} \text{find } & \ddot{q}, \lambda \\ \text{subject to } & H(q)\ddot{q} + C(q,\dot{q}) + G(q) = B(q)u + J(q)^{\mathrm{T}}\lambda \\ & \phi(q) \geq 0 \\ & \lambda \geq 0 \\ & \phi(q)^{\mathrm{T}}\lambda = 0. \end{aligned} \tag{1}$$

where $q \in \mathbb{R}^n$ is a vector of generalized coordinates, $H(\cdot) \in \mathbb{R}^{n\times n}$ is the inertia matrix, $C(\cdot,\cdot) \in \mathbb{R}^n$ are the Coriolis terms, $G(\cdot) \in \mathbb{R}^n$ is the gravitational forces, and $B(\cdot) \in \mathbb{R}^{n\times m}$ is the input mapping of m inputs. $\phi(q) \geq 0$ represents a non-penetration constraint where $\phi(q) : \mathbb{R}^n \to \mathbb{R}^k$ for k potential contacts. $\lambda \in \mathbb{R}^k$ is the vector of constraining normal forces, mapped into generalized coordinates with the Jacobian $J(q)$.

For contact-implicit optimization, the optimization problem is then defined as

$$\underset{\{h, x_0, \ldots, x_N, u_1, \ldots, u_N, \lambda_1, \ldots, \lambda_N\}}{\text{minimize}} g_f(x_N) + h \sum_{k=1}^{N-1} g(x_{k-1}, u_k) \tag{2}$$

where h is the length of the time steps for $k = 1 \ldots N-1$, $g_f(\cdot)$ is the final cost function, and $g(\cdot)$ is the integrated cost function. This optimization is subject to various constraints imposed by the manipulator dynamics as well as the contacts. A key idea is that by optimizing over feasible states x, control inputs u, and contact forces λ at each time step, the optimizer can reason about when and where contacts should optimally occur in a dynamically feasible way. It is important to understand that a major assumption is built into this approach. As presently formulated, the optimizer needs access to the state and dynamics of the entire system, which in the case of this work, consists of the manipulator and the object. Recall that in this work we are not assuming a perfect knowledge of the object beforehand. Instead we rely on an iterative interaction between identification and manipulation to provide the necessary information (Figure 4).

Therefore, I propose the use of contact implicit adaptive model predictive control (CIA-MPC) for planning and control during contact-rich tasks. This controller, because of its adaptive properties borrowed from MRAC, should be robust to the uncertain dynamics of the soft robot arms and therefore be able to take advantage of the benefits of 'mechanical intelligence'.

Contact-implicit trajectory optimization has largely been offline, planar, and has assumed a near-perfect knowledge of the world. Extending it to real-time control of uncertain objects in 3D will allow the robot to simultaneously optimize contact sequences with joint trajectories in a way that scales beneficially with respect to degrees of freedom and number of contacts, without a need for a perfect knowledge of the world beforehand.

To aid in producing real-time capable control, I plan to extend the sampling and parameterization ideas introduced in [83] which enabled real-time MPC for high-DOF systems. An in-depth study on the interactions between object identification and manipulation with CIA-MPC, as well as an algorithm to accomplish this will also be a major contribution of my proposed work.

5 Anticipated Contributions

Apart from my already-published contributions on the dynamic modeling of the soft robot platform [69], [79] which will be used for this research, I anticipate technical contributions in each of the three main objectives listed in Section 3. This research will advance the state-of-the-art capabilities of robotic manipulation and will be communicated with the robotics community via peer-reviewed conference and journal publications. The timeline of these proposed contributions, as well as target publications, are indicated in Figure 5.

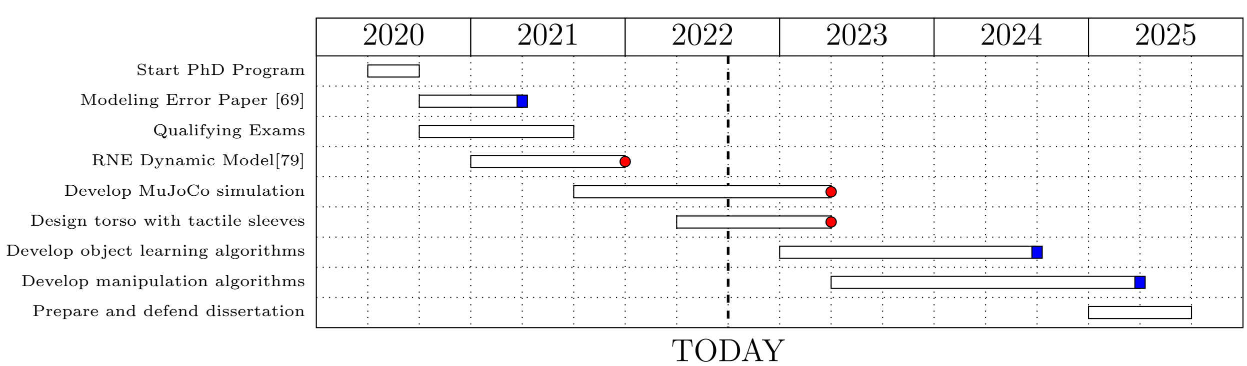

Figure 5: Timeline for accomplishment of proposed research and target conference publications (red dots) and journal publications (blue squares).

References

[1] M. T. Mason, "Toward robotic manipulation," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, p. 1–28, 2018.

[2] R. Tedrake, *Robotic Manipulation*, 2022. [Online]. Available: http://manipulation.mit.edu

[3] K. Salisbury, W. Townsend, B. Ebrman, and D. DiPietro, "Preliminary design of a whole-arm manipulation system (wams)," in *1988 IEEE International Conference on Robotics and Automation Proceedings*, Apr 1988, p. 254–260 vol.1, 245 citations (Semantic Scholar/DOI) [2022-06-23] 136 citations (Crossref) [2022-06-23].

[4] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*, 1st ed. CRC Press, Dec 2017. [Online]. Available: https://www.taylorfrancis.com/books/9781351469791

[5] M. T. Mason, *Mechanics of robotic manipulation*. MIT press, 2001.

[6] O. Brock, J. Kuffner, and J. Xiao, *Motion for Manipulation Tasks*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, p. 615–645. [Online]. Available: https://doi.org/10.1007/978-3-540-30301-5_27

[7] I. Kao, K. Lynch, and J. W. Burdick, *Contact Modeling and Manipulation*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, p. 647–669. [Online]. Available: https://doi.org/10.1007/978-3-540-30301-5_28

[8] D. Prattichizzo and J. C. Trinkle, *Grasping*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, p. 671–700. [Online]. Available: https://doi.org/10.1007/978-3-540-30301-5_29

[9] A. Goncalves, N. Kuppuswamy, A. Beaulieu, A. Uttamchandani, K. M. Tsui, and A. Alspach, "Punyo-1: Soft tactile-sensing upper-body robot for large object manipulation and physical human interaction," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, Apr 2022, p. 844–851.

[10] N. Dehio, Y. Wang, and A. Kheddar, "Dual-arm box grabbing with impact-aware mpc utilizing soft deformable end-effector pads," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, p. 5647–5654, Apr 2022, 0 citations (Semantic Scholar/DOI) [2022-08-18].

[11] P. Gliesche, C. Kowalski, M. Pfingsthorn, A. Hein, and L. Bodenhagen, "Towards whole arm manipulation for outpatient care." Las Vegas, NV: IEEE, Oct 2020, p. 5.

[12] Y. Liu, Z. Ge, S. Yang, I. D. Walker, and Z. Ju, "Elephant's trunk robot: An extremely versatile under-actuated continuum robot driven by a single motor," *Journal of Mechanisms and Robotics*, vol. 11, no. 5, p. 051008, Oct 2019, 12 citations (Semantic Scholar/DOI) [2022-08-23].

[13] N. K. Uppalapati and G. Krishnan, "Design and modeling of soft continuum manipulators using parallel asymmetric combination of fiber-reinforced elastomers," *Journal of Mechanisms and Robotics*, vol. 13, no. 1, p. 011010, Feb 2021, 8 citations (Semantic Scholar/DOI) [2022-08-23].

[14] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. eaat8414, Jun 2019.

[15] B. Zhang, Y. Xie, J. Zhou, K. Wang, and Z. Zhang, "State-of-the-art robotic grippers, grasping and control strategies, as well as their applications in agricultural robots: A review," *Computers and Electronics in Agriculture*, vol. 177, p. 105694, Oct 2020, 68 citations (Semantic Scholar/DOI) [2022-08-23].

[16] W. Ruotolo, D. Brouwer, and M. R. Cutkosky, "From grasping to manipulation with gecko-inspired adhesives on a multifinger gripper," *Science Robotics*, vol. 6, no. 61, p. eabi9773, Dec 2021, 2 citations (Semantic Scholar/DOI) [2022-08-23].

[17] S. Kim, C. Laschi, and B. Trimmer, "Soft robotics: a bioinspired evolution in robotics," *Trends in Biotechnology*, vol. 31, no. 5, p. 287–294, May 2013.

[18] Q. Li, O. Kroemer, Z. Su, F. F. Veiga, M. Kaboli, and H. J. Ritter, "A review of tactile information: Perception and action through touch," *IEEE Transactions on Robotics*, vol. 36, no. 6, p. 1619–1634, Dec 2020, 54 citations (Semantic Scholar/DOI) [2022-07-26].

[19] A. Zhan, P. Zhao, L. Pinto, P. Abbeel, and M. Laskin, "A framework for efficient robotic manipulation," no. arXiv:2012.07975, Dec 2020, 35 citations (Semantic Scholar/arXiv) [2022-08-22] arXiv:2012.07975 [cs]. [Online]. Available: http://arxiv.org/abs/2012.07975

[20] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Scalable deep reinforcement learning for vision-based robotic manipulation," p. 23.

[21] H. Yousef, M. Boukallel, and K. Althoefer, "Tactile sensing for dexterous in-hand manipulation in robotics—a review," *Sensors and Actuators A: Physical*, vol. 167, no. 2, p. 171–187, Jun 2011, 611 citations (Semantic Scholar/DOI) [2022-08-29].

[22] Y. Hou, Z. Jia, and M. T. Mason, "Manipulation with shared grasping," no. arXiv:2006.02996, Jun 2020, 12 citations (Semantic Scholar/arXiv) [2022-06-23] number: arXiv:2006.02996 arXiv:2006.02996 [cs]. [Online]. Available: http://arxiv.org/abs/2006.02996

[23] A. Yamaguchi and C. G. Atkeson, "Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision?" *Advanced Robotics*, vol. 33, no. 14, p. 661–673, Jul 2019, 41 citations (Semantic Scholar/DOI) [2022-08-29].

[24] P. Mittendorfer, E. Yoshida, and G. Cheng, "Realizing whole-body tactile interactions with a self-organizing, multi-modal artificial skin on a humanoid robot," *Advanced Robotics*, vol. 29, no. 1, p. 51–67, Jan 2015, 94 citations (Semantic Scholar/DOI) [2022-07-18].

[25] N. Day, J. Penaloza, V. J. Santos, and M. D. Killpack, "Scalable fabric tactile sensor arrays for soft bodies," *Journal of Micromechanics and Microengineering*, vol. 28, no. 6, p. 064004, Jun 2018, 13 citations (Semantic Scholar/DOI) [2022-07-07].

[26] H. Choi, C. Crump, C. Duriez, A. Elmquist, G. Hager, D. Han, F. Hearl, J. Hodgins, A. Jain, F. Leve, C. Li, F. Meier, D. Negrut, L. Righetti, A. Rodriguez, J. Tan, and J. Trinkle, "On the use of simulation in robotics: Opportunities, challenges, and suggestions for moving forward," *Proceedings of the National Academy of Sciences*, vol. 118, no. 1, p. e1907856118, Jan 2021.

[27] J. Collins, S. Chand, A. Vanderkop, and D. Howard, "A review of physics simulators for robotic applications," *IEEE Access*, vol. 9, p. 51416–51431, 2021.

[28] C. K. Liu and D. Negrut, "The role of physics-based simulators in robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, no. 1, p. 35–58, May 2021, 8 citations (Semantic Scholar/DOI) [2022-08-30].

[29] S. Höfer, K. Bekris, A. Handa, J. C. Gamboa, M. Mozifian, F. Golemo, C. Atkeson, D. Fox, K. Goldberg, J. Leonard, C. Karen Liu, J. Peters, S. Song, P. Welinder, and M. White, "Sim2real in robotics and automation: Applications and challenges," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 2, p. 398–400, Apr 2021, 9 citations (Semantic Scholar/DOI) [2022-08-30].

[30] P. C. Horak and J. C. Trinkle, "On the similarities and differences among contact models in robot simulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, p. 493–499, Apr 2019.

[31] R. Elandt, E. Drumwright, M. Sherman, and A. Ruina, "A pressure field model for fast, robust approximation of net contact force and moment between nominally rigid objects," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2019, p. 8238–8245, 11 citations (Semantic Scholar/DOI) [2022-08-11].

[32] M. Körber, J. Lange, S. Rediske, S. Steinmann, and R. Glück, "Comparing popular simulation environments in the scope of robotics and reinforcement learning," no. arXiv:2103.04616, Mar 2021, 5 citations (Semantic Scholar/arXiv) [2022-08-30] arXiv:2103.04616 [cs]. [Online]. Available: http://arxiv.org/abs/2103.04616

[33] T. Erez, Y. Tassa, and E. Todorov, "Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ode and physx," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, p. 4397–4404, 194 citations (Semantic Scholar/DOI) [2022-08-30].

[34] B. S. Eberman, “Whole-arm manipulation: Kinematics and control,” Master’s thesis, MIT, Jan 1989. [Online]. Available: https://dspace.mit.edu/bitstream/handle/1721.1/14428/20242325-MIT.pdf?sequence=2

[35] P. Song, M. Yashima, and V. Kumar, “Dynamics and control of whole arm grasps,” in *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, vol. 3, May 2001, p. 2229–2234 vol.3, 25 citations (Semantic Scholar/DOI) [2022-06-23] 12 citations (Crossref) [2022-06-23] ISSN: 1050-4729.

[36] F. Asano, Y. Saitoh, K. Watanabe, Z.-W. Luo, and M. Yamakita, “On dynamic whole body manipulation,” in *Proceedings 2003 IEEE International Symposium on Computational Intelligence in Robotics and Automation. Computational Intelligence in Robotics and Automation for the New Millennium (Cat. No.03EX694)*, vol. 3, Jul 2003, p. 1201–1206 vol.3, 7 citations (Semantic Scholar/DOI) [2022-06-23].

[37] D. Devereux, “Control strategies for whole arm grasping,” Ph.D. dissertation, University of Manchester, 2010.

[38] A. Jain, M. D. Killpack, A. Edsinger, and C. C. Kemp, “Reaching in clutter with whole-arm tactile sensing,” *The International Journal of Robotics Research*, vol. 32, no. 4, p. 458–482, Apr 2013, 98 citations (Semantic Scholar/DOI) [2022-06-23].

[39] M. D. Killpack and C. C. Kemp, “Fast reaching in clutter while regulating forces using model predictive control,” in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Oct 2013, p. 146–153, 15 citations (Semantic Scholar/DOI) [2022-10-24].

[40] M. Florek-Jasińska, T. Wimböck, and C. Ott, “Humanoid compliant whole arm dexterous manipulation: Control design and experiments,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep 2014, p. 1616–1621, 10 citations (Semantic Scholar/DOI) [2022-06-23] 6 citations (Crossref) [2022-06-23] ISSN: 2153-0866.

[41] M. Neunert, M. Stäuble, M. Giftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli, “Whole-body nonlinear model predictive control through contacts for quadrupeds,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, p. 1458–1465, Jul 2018, 120 citations (Semantic Scholar/DOI) [2022-07-26].

[42] T. Marcucci and R. Tedrake, “Warm start of mixed-integer programs for model predictive control of hybrid systems,” *IEEE Transactions on Automatic Control*, vol. 66, no. 6, p. 2433–2448, Jun 2021, 22 citations (Semantic Scholar/DOI) [2022-07-26].

[43] A. Anders, “Learning a strategy for whole-arm grasping,” Master’s thesis, 2014.

[44] A. Murali, Y. Li, D. Gandhi, and A. Gupta, “Learning to grasp without seeing,” no. arXiv:1805.04201, May 2018, 37 citations (Semantic Scholar/arXiv) [2022-07-26] arXiv:1805.04201 [cs]. [Online]. Available: http://arxiv.org/abs/1805.04201

[45] A. Colomé and C. Torras, *Reinforcement Learning of Bimanual Robot Skills*, ser. Springer Tracts in Advanced Robotics. Cham: Springer International Publishing, 2020, vol. 134. [Online]. Available: http://link.springer.com/10.1007/978-3-030-26326-3

[46] M. Saveriano, F. J. Abu-Dakka, A. Kramberger, and L. Peternel, "Dynamic movement primitives in robotics: A tutorial survey," no. arXiv:2102.03861, Feb 2021, 13 citations (Semantic Scholar/arXiv) [2022-09-13] arXiv:2102.03861 [cs]. [Online]. Available: http://arxiv.org/abs/2102.03861

[47] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, "Tactile dexterity: Manipulation primitives with tactile feedback," no. arXiv:2002.03236, Apr 2020, 30 citations (Semantic Scholar/arXiv) [2022-06-27] number: arXiv:2002.03236 arXiv:2002.03236 [cs, eess]. [Online]. Available: http://arxiv.org/abs/2002.03236

[48] N. Vuong, H. Pham, and Q.-C. Pham, "Learning sequences of manipulation primitives for robotic assembly," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, p. 4086–4092, 5 citations (Semantic Scholar/DOI) [2022-09-13].

[49] W. Wan, C. Sun, J. Yuan, X. Hou, Y. Guo, Y. Ou-yang, Q. Li, L. Zhao, H. Shi, and D. Han, "Adaptive whole-arm grasping approach of tumbling space debris by two coordinated hyper-redundant manipulators," in *Intelligent Robotics and Applications*, ser. Lecture Notes in Computer Science, H. Yu, J. Liu, L. Liu, Z. Ju, Y. Liu, and D. Zhou, Eds. Cham: Springer International Publishing, 2019, p. 450–461.

[50] M. Murooka, Y. Inagaki, R. Ueda, S. Nozawa, Y. Kakiuchi, K. Okada, and M. Inaba, "Whole-body holding manipulation by humanoid robot based on transition graph of object motion and contact," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2015, p. 3950–3955, 4 citations (Semantic Scholar/DOI) [2022-07-12].

[51] J. Carius, R. Ranftl, V. Koltun, and M. Hutter, "Trajectory optimization with implicit hard contacts," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, p. 3316–3323, Oct 2018, 38 citations (Semantic Scholar/DOI) [2022-07-26].

[52] V. Kurtz and H. Lin, "Contact-implicit trajectory optimization with hydroelastic contact and ilqr," no. arXiv:2202.13986, Feb 2022. [Online]. Available: http://arxiv.org/abs/2202.13986

[53] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research*, vol. 33, no. 1, p. 69–81, Jan 2014, 420 citations (Semantic Scholar/DOI) [2022-07-26].

[54] A. Patel, S. L. Shield, S. Kazi, A. M. Johnson, and L. T. Biegler, "Contact-implicit trajectory optimization using orthogonal collocation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, p. 2242–2249, Apr 2019, 31 citations (Semantic Scholar/DOI) [2022-07-26].

[55] Z. Manchester and S. Kuindersma, *Variational Contact-Implicit Trajectory Optimization*, ser. Springer Proceedings in Advanced Robotics. Cham: Springer International Publishing, 2020, vol. 10, p. 985–1000.

[56] C. Wang, X. Zhang, X. Zang, Y. Liu, G. Ding, W. Yin, and J. Zhao, "Feature sensing and robotic grasping of objects with uncertain information: A review," *Sensors*, vol. 20, no. 13, p. 3707, Jul 2020, 15 citations (Semantic Scholar/DOI) [2022-07-26].

[57] N. Mavrakis and R. Stolkin, "Estimation and exploitation of objects ' inertial parameters in robotic grasping and manipulation: A survey," *Robotics and Autonomous Systems*, vol. 124, p. 103374, Feb 2020, 18 citations (Semantic Scholar/DOI) [2022-09-26].

[58] W. Gao and R. Tedrake, "kpam-sc: Generalizable manipulation planning using keypoint affordance and shape completion," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, p. 6527–6533, 25 citations (Semantic Scholar/DOI) [2022-07-26].

[59] X. Zhang, Z. Zhang, C. Zhang, J. Tenenbaum, B. Freeman, and J. Wu, "Learning to reconstruct shapes from unseen classes," in *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc., 2018. [Online]. Available: https://proceedings.neurips.cc/paper/2018/hash/208e43f0e45c4c78cafadb83d2888cb6-Abstract.html

[60] W. Yuan, S. Dong, and E. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, Nov 2017, 300 citations (Semantic Scholar/DOI) [2022-10-05].

[61] N. Kuppuswamy, A. Alspach, A. Uttamchandani, S. Creasey, T. Ikeda, and R. Tedrake, "Soft-bubble grippers for robust and perceptive manipulation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2020, p. 9917–9924, 29 citations (Semantic Scholar/DOI) [2022-10-05].

[62] Y. Yu, K. Fukuda, and S. Tsujio, "Estimation of mass and center of mass of graspless and shape-unknown object," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 4, May 1999, p. 2893–2898 vol.4, 30 citations (Semantic Scholar/DOI) [2022-09-26].

[63] C. Wang, S. Wang, B. Romero, F. Veiga, and E. Adelson, "Swingbot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2020, p. 5633–5640, 26 citations (Semantic Scholar/DOI) [2022-07-26].

[64] M. Murooka, S. Nozawa, Y. Kakiuchi, K. Okada, and M. Inaba, "Feasibility evaluation of object manipulation by a humanoid robot based on recursive estimation of the object's physical properties," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, May 2017, p. 4082–4089, 5 citations (Semantic Scholar/DOI) [2022-09-28].

[65] S. Li, S. A. Awale, K. E. Bacher, T. J. Buchner, C. Della Santina, R. J. Wood, and D. Rus, "Scaling up soft robotics: A meter-scale, modular, and reconfigurable soft robotic system," *Soft Robotics*, vol. 9, no. 2, p. 324–336, Apr 2022.

[66] N. S. Usevitch, Z. M. Hammond, M. Schwager, A. M. Okamura, E. W. Hawkes, and S. Follmer, "An untethered isoperimetric soft robot," *Science Robotics*, vol. 5, no. 40, p. eaaz0492, Mar 2020, 40 citations (Semantic Scholar/DOI) [2022-10-08].

[67] H. Mao and J. Xiao, *Object Shape Estimation Through Touch-Based Continuum Manipulation*, ser. Springer Proceedings in Advanced Robotics. Cham: Springer International Publishing, 2020, vol. 10, p. 573–588. [Online]. Available: http://link.springer.com/10.1007/978-3-030-28619-4_42

[68] A. D. Marchese, R. K. Katzschmann, and D. Rus, "Whole arm planning for a soft and highly compliant 2d robotic manipulator," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep 2014, p. 554–560, 94 citations (Semantic Scholar/DOI) [2022-06-23].

[69] C. C. Johnson, T. Quackenbush, T. Sorensen, D. Wingate, and M. D. Killpack, "Using first principles for deep learning and model-based control of soft robots," *Frontiers in Robotics and AI*, vol. 8, 2021. [Online]. Available: https://www.frontiersin.org/articles/10.3389/frobt.2021.654398/full

[70] T. G. Thuruthel, E. Falotico, F. Renda, and C. Laschi, "Model-based reinforcement learning for closed-loop dynamic control of soft robotic manipulators," *IEEE Transactions on Robotics*, vol. 35, no. 1, p. 124–134, Feb 2019.

[71] M. A. Graule, C. B. Teeple, T. P. McCarthy, G. R. Kim, R. C. St. Louis, and R. J. Wood, "Somo: Fast and accurate simulations of continuum robots in complex environments," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2021, p. 3934–3941, 4 citations (Semantic Scholar/DOI) [2022-07-07].

[72] Y. Toshimitsu, K. W. Wong, T. Buchner, and R. Katzschmann, "Sopra: Fabrication and dynamical modeling of a scalable soft continuum robotic arm with integrated proprioceptive sensing," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2021, p. 653–660.

[73] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: https://drake.mit.edu

[74] R. K. Katzschmann, C. D. Santina, Y. Toshimitsu, A. Bicchi, and D. Rus, "Dynamic motion control of multi-segment soft robots using piecewise constant curvature matched with an augmented rigid body model," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, Apr 2019, p. 454–461, 50 citations (Semantic Scholar/DOI) [2022-08-11].

[75] Y. Hu, J. Liu, A. Spielberg, J. B. Tenenbaum, W. T. Freeman, J. Wu, D. Rus, and W. Matusik, "Chainqueen: A real-time differentiable physical simulator for soft

robotics," in *2019 International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE, May 2019, p. 6265–6271. [Online]. Available: https://ieeexplore.ieee.org/document/8794333/

[76] A. Spielberg, T. Du, Y. Hu, D. Rus, and W. Matusik, "Advanced soft robot modeling in chainqueen," *Robotica*, p. 1–31, Jun 2021, 0 citations (Semantic Scholar/DOI) [2022-08-11].

[77] E. Coevoet, T. Morales-Bieze, F. Largilliere, Z. Zhang, M. Thieffry, M. Sanz-Lopez, B. Carrez, D. Marchal, O. Goury, J. Dequidt, and C. Duriez, "Software toolkit for modeling, simulation, and control of soft robots," *Advanced Robotics*, vol. 31, no. 22, p. 1208–1224, Nov 2017.

[78] P. Hyatt, C. C. Johnson, and M. D. Killpack, "Model reference predictive adaptive control for large-scale soft robots," *Frontiers in Robotics and AI*, vol. 7, 2020. [Online]. Available: https://www.frontiersin.org/articles/10.3389/frobt.2020.558027/full

[79] S. W. Jensen, C. C. Johnson, A. M. Lindberg, and M. D. Killpack, "Tractable and intuitive dynamic model for soft robots via the recursive newton-euler algorithm," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, Apr 2022, p. 416–422.

[80] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2012, p. 5026–5033, 3048 citations (Semantic Scholar/DOI) [2022-10-08].

[81] L. Howell, P. A. Midha, and T. W. Norton, "Evaluation of equivalent spring stiffness for use in a pseudo-rigid-body model of large-deflection compliant mechanisms."," *Journal of Mechanical Design*, p. 126–131, 1996.

[82] J. Varley, C. DeChant, A. Richardson, J. Ruales, and P. Allen, "Shape completion enabled robotic grasping," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2017, p. 2442–2447, 190 citations (Semantic Scholar/DOI) [2022-09-26].

[83] P. E. Hyatt, "Robust real-time model predictive control for high degree of freedom soft robots," p. 232.