

Optimal Kinodynamic Planning for Compliant Mobile Manipulators

Katie Byl

Abstract— Robot manipulation tasks for a mobile, personal robot will often be importantly distinct from those of a traditional, factory robot arm; correspondingly, appropriate motion planning solutions may be notably different, as well. This paper introduces novel definitions of and solution methods for optimal kinodynamic planning for mobile robot manipulation applications. In particular, we consider a generalized, real-world scenario where (1) robots must interact safely with humans and with objects in home environments, (2) there may be significant uncertainty about the impedance properties of objects to be manipulated, and (3) there may also be significant environmental and sensory noise. Unlike traditional kinodynamic planning, where minimum time trajectories are generally considered optimal, we suggest that the paramount goal in planning for personal robots should instead be safety; that is, we should minimize the probability of “failure” or at least ensure it is below some threshold. We discuss the appropriateness of this definition of optimality, suggest a generalized methodology for achieving optimal or near-optimal solutions, and present a simple, 2D manipulation problem to demonstrate the approach.

I. INTRODUCTION

This paper discusses the problem of selecting a particular initial robot pose and subsequent manipulator trajectories for a mobile robot with arms. This problem is a particular case of kinodynamic planning: finding motions for a robot that obey both (1) kinematic constraints, e.g., positions with feasible kinematic configurations that avoid collisions with obstacles, and (2) dynamic (or *differential*) constraints, e.g. velocities and accelerations that do not exceed actuator limitations. This general kinodynamic problem becomes more challenging for robots that operate in environments where safety is paramount and where the environment is not perfectly known, and we believe the problem statement for kinodynamic planning in this regime should correspondingly be modified.

For kinodynamic motion planning, optimality typically means a *minimum time* solution – getting from a start to an end configuration as quickly as possible – which is an appropriate definition for traditional, high-impedance robot arms. We suggest a new definition of *optimality* is required when planning motions for personal robots, and we present a toy example to illustrate a proposed method for deriving corresponding optimal

plans. Specifically, we propose that the paramount goal for mobile manipulators should be to minimize the chance of “failure”, where failure is defined to encompass all undesirable outcomes, as discussed further in Section II. Employing safer, friendlier robots (appropriate for the home) is a practical necessity for the future, but making robots safer is likely to increase the impact that noise and uncertainty have on the dynamics of robot manipulation tasks. There is a need for improved motion planning techniques that can derive optimal control policies in this important, new regime.

A. Low-impedance personal robots

Mechanical impedance, which is analogous to electrical impedance, is the frequency-dependent relationship describing the force required to achieve a particular velocity of motion:

$$Z(s) = \frac{F(s)}{V(s)} \quad (1)$$

Traditional, factory-use robot manipulators are designed with speed and accuracy as obvious goals and correspondingly have quite high mechanical impedance (i.e., stiffness), while humans and other animals interact with the world with a much higher degree of compliance [24]. A factory or laboratory environment for a robot is generally well-characterized, with little uncertainty; in this regime, fast and powerful robotics are a practical solution. By contrast, in situations where there is significant uncertainty, due to incomplete or noisy sensing, both about geometry of the environment and about the impedance (i.e., mass, viscous damping and/or friction, spring stiffness) of the objects to be manipulated, a stiff robot is simply not a safe or practical solution: it may seriously damage itself, the environment and nearby humans during both unplanned collisions and precisely-planned manipulation tasks. It is correspondingly becoming increasingly clear that robots designed with non-trivial compliance are a practical solution if robots and robotic devices are to leave well-characterized factory and laboratory settings and enter our homes [25], [23], [3], [10], [31], [33], [15].

While there are important motivations for designing personal robots with low-impedance: for safety and comfort, and to exploit the “forgiving” effects that compliant passive dynamics may lend during inter-

This work was not supported by any organization
K. Byl is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106, USA (katiebyl@ece.ucsb.edu).

actions between the robot and an uncertain environment, there are also trade-offs in adopting such a design. Lower robot stiffness inherently means less ability to resist perturbations, resulting in lower positional accuracy. Using lower forces to accomplish manipulation tasks will also generally result in lower bandwidth [13], and the dynamic behavior of a low-impedance robot arm is coupled, nonlinear and depends highly on the arm’s configuration in space [1].

B. Kinodynamic Planning

Over the past 25 years, a substantial body of work has developed on kinodynamic motion planning for robots. Initial approaches studied trajectory planning when a particular path is specified [4], with inertial constraints [22], and over a discrete, tessellated subspace of the full problem [26]. This was followed by significant results in optimal and near-optimal solutions for low-dimensional, classical kinodynamic planning [9], [8], [12], [11].

For higher degree-of-freedom robots, however, finding optimal solutions becomes increasingly challenging, and many popular strategies used today involve sampling state space to find feasible but non-optimal solutions. Such approaches include rapidly-exploring randomized trees (RRTs) [20], [19], [21], including cases in which obstacles move over time [14]; probabilistic roadmaps [18]; and a variety of other, sample-based approaches [2], [32], [28], [16]. These approaches generally involve sequential planning, where a feasible kinematic plan is found first and the “speed” with which it is executed is then adjusted to ensure dynamic feasibility. For underactuated, dynamic robot motions (e.g., locomotion and manipulation), such sequential kinodynamic planning methods may no longer be possible, since motions are coupled and will only succeed when performed at appropriate speed [6], [27]. The development of new motion planning techniques for such underactuated systems [30] is an area of active research, as are methods for insuring that sample-based planners converge upon optimal solutions [17]. Developing practical methods to generate optimal motion control strategies for underactuated and often underpowered tasks such as manipulation and locomotion in an important and active field of research; this paper suggests a new modeling and design approach toward this ultimate goal.

II. PROBLEM STATEMENT

Central to this paper, we argue that in compliant, mobile manipulation, it may often be impossible to guarantee the successful completion of a desired task, no matter how conservatively we design a motion plan. In a real-world environment, where the geometry and impedance characteristics of the surroundings are not known perfectly, it may be impossible to know

with absolute certainty that a particular manipulation task will not result in *failure*. Examples of failures include: collisions, infeasible kinematics, dropping a poorly-grasped object, insufficient actuator power, excessive time spent accomplishing a task, damage to a person or object during a planned interaction, and/or any other task-specific outcome defined as unacceptable.

All hope is not lost, however! Humans and other animals are subject to the same difficulties, and they can often perform with astonishing agility and dexterity, nonetheless. Even without absolute guarantees of success, one can still plan motions that make the risk of failure acceptably (or even vanishingly) low. We propose the following definition of “optimality” for the problem of kinodynamic motion planning for mobile manipulators:

DEFINITION: Optimal kinodynamic motion plans are those that *minimize the probability of failure*.

In Section III, we outline a method to find an optimal solution, using a variation of an approach used in finding near-optimal solutions for an underactuated biped model on rough terrain [7], [5].

III. GENERAL THEORY

This section outlines a general framework for designing reliable robot motion control, in which the robot and its environment are together modeled as a Markov decision process (MDP). We focus first on the rationale for this approach, without outlining an exact approach, which may depend heavily on the types of uncertainty there are in a particular task. This is followed by a toy example in Section IV.

Inherent in our modeling throughout is the assumption of a “go/no-go” outcome for a robot task. In legged locomotion, for example, one obvious no-go status is “fallen down”, and the dynamics of manipulation tasks are similarly punctuated by dramatic, discrete events: objects can fall out of grasp, spring-loaded doors may resist opening, and collisions can occur. One obvious objective is: *to maximize the likelihood a robot will successfully complete one or more tasks*. We emphasize here that this is usually the true primary goal in motion planning for a robot. Minimizing some quantifiable combination of undesirable costs, such as time or energy, is frequently a secondary goal.

For particular robot manipulation tasks in well-characterized environments, it may be both quite valid and practical to assume that success can be guaranteed. For example, this may be the case when the environment is well-characterized, sensors have little noise, and we are confident of achieving sufficient actuator torques and a good grasp wrench for a given set of tasks. In this regime, we take success for granted

and can instead focus on optimizing speed or power consumption.

By contrast, bringing robots out of structured environments and into often cluttered (and changing) home environments generally makes such “near-guarantees” of success much more unlikely, as do any reductions in actuator torque outputs that are required to ensure safety. In this regime, we suggest the following general approach:

1. Carefully define “success” for a given task (or set of tasks).
2. Estimate the critical sources of uncertainty.
3. Create a reduced-order Markov model of the dynamic system.
4. Solve for the control policy that maximizes the probability of success.

For step 1, success will often require moving an object “from A to B”. However, the definition of success can also incorporate notions of time and/or power limits for completion of a task, to ensure that the robot does not move with seemingly unreasonable trepidation. We can also require that motions do not startle humans nearby, or that they leave explicit safety margins for collisions (in particular, around animate objects).

In estimating uncertainty (step 2), it may often be challenging and impractical to quantify exact probability distributions. Fortunately, identifying *what* is uncertain – in the geometry and/or impedance of system, for instance – is often more critical than precise quantification of degree. For example, modeling the mass of an object as uncertain may automatically result motion plans that ensure a robot arm is oriented in a way that best exploits the non-linear relationship between manipulator geometry and force output at the end effector. When possible, however, one should avoid being unnecessarily conservative about less significant sources of uncertainty, as this may result in awkward and wasteful motions that are not actually beneficial. Developing improved techniques for estimating the sources and magnitudes of uncertainty in a robot manipulation task is an important problem both for the success of this approach and in personal robotics, generally.

Capturing the essential dynamics for step 3 could theoretically involve an enormous number of potential states at times, if we have many sources of uncertainty across several degrees of freedom. Thankfully, many of the unknowns in geometry and impedance in a system will remain constant throughout a particular trial, so the transition matrix describing the dynamics

will be quite sparse, simplifying analysis or even allowing us to break the problem into a set of “what if” scenarios that are each straight-forward to analyze (maintaining several, smaller transition matrices that cannot cross-communicate, or even without creating any transition matrix at all).

Finally, once we have described the desired task(s) of a stochastic dynamic system, we can employ techniques in machine learning to solve for an optimal policy (step 4). When good models are practical, this step may simply involve policy iteration: alternately solving Bellman’s optimality equation (to find the value, $V_\pi(s)$, of every state for a given policy) and calculating the best policy ($\pi(s)$) for a given definition of $V_\pi(s)$, until this iterative process converges to an optimal policy; and when such models are not practical, we can employ algorithms for learning online [29].

IV. TOY EXAMPLE

In this section, we: (1) present a toy model of a mobile robot manipulator in an uncertain (but characterizable) environment, (2) define a desired set of tasks and the dominant failure modes that may occur, and (3) solve for the optimal kinodynamic motion plan that minimizes the probability of failure.




Fig. 1. 2D robot manipulator model

Figure 1 shows a simple, 2D manipulator model. We assume the robot can move to a desired position, x_p , in preparation for a manipulation task (where it remains throughout the task), and that it has a telescoping torso that brings the shoulder to some height, z ; both x_p and z are measured with respect to a global coordinate system with origin (x_o, z_o) . The arm itself has only two degrees of freedom: θ_1 and θ_2 , which define the orientations of the first two links (L_1 and L_2). For this toy model, we assume the third link is constrained (through control) to remain parallel to the ground. Given the variables x_p , z , θ_1 , and θ_2 , we assume (for simplicity) that contact with an object occurs at a particular “end effector” location (x_e, z_e) , at a length L_3 along the hand, so:

$$x_e = x_p + L_1 \sin \theta_1 + L_2 \sin (\theta_1 + \theta_2) + L_3 \quad (2)$$

$$z_e = z - L_1 \cos \theta_1 - L_2 \cos (\theta_1 + \theta_2) \quad (3)$$

Additionally, we will assume that each joint has a torque limit (2 N-m). The range of possible forces at the end effector then depend on the impedance properties of the robot arms (mass distribution, etc.), and the current state (kinematic configuration and joint velocities) of the robot. (We will not detail the exact model assumed here, for brevity.)

As a first step, we define a set of desired manipulation tasks that must be successfully completed to achieve. Referencing the cartoon in Figure 2, Task 1 is to move a flower vase of unknown mass, m_v , from the top shelf to the lower shelf and back again (e.g., to clean or inspect the shelf), and Task 2 is to fold down the lower shelf. This lower shelf has some mass distribution, m_s , and it is connected to a fixed wall via a linear spring, k_s , and a damper b_s . This mass, stiffness and damping constitute some generalized mechanical impedance, and in this problem, we will assume these properties are known very well. To focus on basic ideas more clearly in this paper, we will also require that both tasks be done relatively slowly. This allows us to use a quasi-static force balance (e.g., that \dot{x}_p , \ddot{x}_p , \dot{z} and \ddot{z} are all very small) for our analysis. In general, however, such dynamics would be not be ignored: they would simply and elegantly be incorporated in the full MDP model that defines the possible state-to-state transitions.

If all of the impedances contributing to the quasi-static forces (i.e., m_s , k_s , and m_v) are small enough, and if we assume perfect knowledge of the geometry and a perfect ability to grasp, then this problem reduces to one of kinematics, only; we just need to ensure the robot is in an appropriate location (x_p) where the end effector can reach all planned motion trajectories. Figure 3 shows an example of the reachable workspace for a particular x_p . The two areas bounded by dashed lines show regions when z is fixed at either a minimum or maximum value; the solid boundary defines the total reachable workspace if the robot is allowed to extend its torso during manipulation.

Of course, we are specifically interested in cases where the problem involves more than kinematics, so (as step 2), let us now identify any critical sources of uncertainty. As an example, let us first assume we are uncertain if the vase is full of water or not. Figure 4 shows the (somewhat arbitrary) probability density function we will use to quantify our uncertainty about the mass of the vase, m_v , in our toy example. Additionally, we will assume that we are very familiar with the impedance properties of the shelf (m_s and k_s




Fig. 2. Example tasks: manipulating objects with impedance.

matter, but not b_s , given our quasi-static assumption), and we will also assume for now (perhaps somewhat naively) that the geometry can be perfectly sensed and controlled.

Figure 5 shows the allowable locations, (x_p, z) , for the robot such that it is kinematically and/or dynamically capable of various subsets of the required tasks. The largest area (in gray, blue and red) is the region in which grasping the vase is kinematically feasible for the manipulator; the next smallest (blue and red) is the subset of this region for which lowering the shelf (assuming a particular contact point on the shelf) is also possible; and the smallest region (red, only) shows the subset of *these* robot positions at which the robot can apply adequate torque at each moment as it slowly lowers the shelf. This final region would also correspond to robot stance locations where we could theoretically guarantee both tasks will be completed with probability one *if we could determine that m_v if was appropriately small enough.*

Now (as step 3), we can consider the effect of our uncertainty in m_v on the dynamics of success and failure in completing our tasks. For this simple example, we need only consider the probability that the weight of the vase will exceed the vertical force the robot can generate at the end effector *at any point along the trajectory.* Here, we analyze only the shortest path for the vase: two straight line segments,




Fig. 3. “Reachable” workspace, for fixed x_p and full range of Δz




Fig. 4. Probability density function (PDF) for mass of vase

first to the left (off of the top shelf) and then down (to the lower shelf). For any given robot location, we can solve for the lowest (“worst case”) threshold in achievable F_z while following the desired end effector path; the probability of failure is simply the area in Figure 4 above the corresponding critical mass, F_z/g . Figure 6 presents a contour plot of the resulting probability of success across all (x_p, z) combinations (assuming z remains constant throughout both tasks); kinematically infeasible regions have a probability that is strictly zero (shown in dark blue), of course.

The optimal (x_p, z) location on Figure 6 is labeled with a large, green dot. For our simple example, this 2D robot pose location is the only control decision (step 4) to be made. Notice that this point is quite




Fig. 5. Feasible robot locations for deterministic tasks




Fig. 6. Probability of success, for exact kinematics and PDF of m_v

close to the edge of the boundary of the kinematically feasible region; this is because we assumed geometry was known and controlled perfectly in our model. Intuitively, one would expect the best solution would keep the robot within some “safety margin” of feasibility. We have intentionally provided a toy example that violates this to emphasize our point throughout: that the real goal is to maximize success. A kinematic safety margin inherently implies there is additional uncertainty somewhere in the system, and a margin for error will naturally emerge if we appropriately include geometric uncertainty in the underlying model. As mentioned in Section III, even a rough (conservative) estimate of uncertainty is often adequate to capture the dominant routes to failure.

To illustrate the inclusion of kinematic uncertainty in our discretized representation of the dynamic system, we will now assume a probability mass distribution, shown in Figure 7, for errors in the initial position of the vase. This PMF corresponds to a standard deviation of 1 cm in the actual position of the vase, and we assume that this error can be detected as the arm approaches the vase (so it would only cause failure if the reach was infeasible; not because the vase is ultimately grasped in the wrong location).




Fig. 7. Uncertainty in geometry (e.g., x_v) can also be modeled

Propagating the effects of this new noise in the geometric location of the vase, we obtain a revised

estimate of the overall probability of success as a function of the robot stance location, shown in both Figures 8 and 9. (Only the color axes differentiate the two plots; Figure 9 shows any region with a probability less than 80% as dark blue; only regions with zero probability are that dark in Figure 8.) The optimal robot pose location has now shifted further from the edge of the “no-noise” feasible kinematic region boundary, which is overlaid as a dashed boundary in both Figures 8 and 9, for better comparison with Figure 6.




Fig. 8. Probability of success, both m_v and x_v are uncertain




Fig. 9. A rescaled color axis (see text), to highlight danger zones

This toy example is primarily intended to motivate the central idea of maximizing the probability of success. Additionally, it illustrates how the familiar concept of a safety margin emerges naturally as a consequence of reasonable, approximate models of the uncertainty. Even arbitrarily-defined safety margins may sometimes be adequate, as they often act indirectly as a means of increasing our probability of success. However, the kinematics and dynamics of manipulation problems are highly nonlinear: we can do much better than an arbitrary margin, and it is certainly possible for inappropriate safety margins to lower our probability of success instead of improving it ¹!

¹The same statements are true of stability margins for walking robots, such as the zero-moment point (ZMP) margin: such margins are really an indirect tool toward our true goal of “success”, and we should be able to do better by modeling the underlying dynamics.

V. FUTURE WORK

Briefly but importantly, we note that although the *policy* in our toy example consisted only of selecting a static pose for the robot, in general, one would develop a rich policy for controlling the various positions, velocities and forces of the robot to accomplish a task. Methods for doing so are the subject of future work.

VI. CONCLUSIONS AND DISCUSSION

A. Conclusions

We propose a new approach for kinodynamic planning for mobile robots in real-world environments, defining optimal motion plans as those that maximize the probability of success at a desired set of tasks. Real-world environments often present inherent, non-trivial stochasticity, and compliant robots may be particularly sensitive to environmental perturbations. As a consequence, it becomes increasingly relevant for an autonomous robot to select actions that will directly minimize the probability that these uncertainties will result in failure.

We further propose that one method for minimizing failure rate is to model the robot, its environment and any desired task(s) as a Markov decision process (MDP) and to solve for the control policy that minimizes the probability of transitioning into an absorbing, failure state. Specifically, we wish to minimize the long-term average rate of failure events, which can be determined from the transition matrix [7]. In this paper, we solved for a trivial “pose” policy, testing discrete combinations of body poses to find an approximate optimal solution; a fuller policy can be found by incorporating more dynamics (velocities and accelerations), by allowing a fuller set of end effector trajectories, and by using gradient-based learning methods to converge on a locally-optimal policy.

B. Discussion

A key step in modeling the system dynamics is to characterize the level of uncertainty there is in the environment. This can be a challenging problem in itself, and we suggest it is therefore an important avenue for future research in robotics. However, we expect that even bad (e.g., highly conservative) estimates of the variability of the system will often provide reasonable (near-optimal) solutions. Identifying the bounds in error for which this statement is true is yet another direction for research.

Although we focus on robot manipulation in this paper, the same general ideas are applicable toward any robotic or mechatronic system where failures may occur due to underactuation and/or under-powered actuators, or due to environmental uncertainty and/or stochasticity. Pretending that failures can be avoided absolutely is naive; it is more practical and helpful

toward the advancement experimental robotics to address the critical issues of estimating and minimizing failure rates directly.

REFERENCES

- [1] Haruhiko Asada and Kamal Youcef-Toumi. *Direct-Drive Robots - Theory and Practice*. The MIT Press, 1987.
- [2] Barraquand, J., Langlois, B., Latombe, and J.-C. Numerical potential field techniques for robot path planning. *Advanced Robotics, 1991. 'Robots in Unstructured Environments', 91 ICAR., Fifth International Conference on*, pages 1012–1017 vol.2, Jun 1991.
- [3] Joaquin Blaya and Hugh Herr. Adaptive control of a variable-impedance ankle-foot orthosis to assist drop-foot gait. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 12(1):24–31, Mar 2004.
- [4] J.E. Bobrow, S. Dubowsky, and J.S. Gibson. Time-optimal control of robotic manipulators along specified paths. *Int. J of Robotics Research*, 4(3):3–17, 1985.
- [5] Katie Byl. *Metastable Legged-Robot Locomotion*. PhD thesis, MIT, Sep 2008.
- [6] Katie Byl and Russ Tedrake. Dynamically diverse legged locomotion for rough terrain. In *Proceedings of the IEEE/RAS International Conference on Robotics and Automation (ICRA), video submission*, May 2009.
- [7] Katie Byl and Russ Tedrake. Metastable walking machines. *International Journal of Robotics Research*, 28(8):1040–1064, August 1 2009.
- [8] John Canny, Ashutosh Rege, and John Reif. An exact algorithm for kinodynamic planning in the plane. In *SCG '90: Proceedings of the sixth annual symposium on Computational geometry*, pages 271–280. ACM, 1990.
- [9] John F. Canny, Bruce Randall Donald, John H. Reif, and Patrick G. Xavier. On the complexity of kinodynamic planning. In *29th Annual Symposium on Foundations of Computer Science*, pages 306–316. IEEE, 1988.
- [10] Aaron M. Dollar and Robert D. Howe. Designing robust robotic graspers for unstructured environments. *Workshop on Manipulation for Human Environments*, Aug. 19, 2006.
- [11] B.R. Donald and P.G. Xavier. Provably good approximation algorithms for optimal kinodynamic planning for cartesian robots and open-chain manipulators. *Algorithmica*, 14(6):480–530, 1995.
- [12] Bruce Donald, Patrick Xavier, John Canny, and John Reif. Kinodynamic motion planning. *J. of the ACM*, 40(5):1048–1066, 1993.
- [13] Steven D. Eppinger and Warren P. Seering. Understanding bandwidth limitations in robot force control. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 904–909, 1987.
- [14] David Hsu, Robert Kindel, Jean-Claude Latombe, and Stephen Rock. Randomized kinodynamic motion planning with moving obstacles. *The International Journal of Robotics Research*, 21(3):233–255, 2002.
- [15] Sang-Ho Hyon. Compliant terrain adaptation for biped humanoids without measuring ground surface and contact forces. *IEEE Transactions on Robotics*, 25(1):171–178, Feb. 2009.
- [16] Maciej Kalisiak. *Toward More Efficient Motion Planning with Differential Constraints*. PhD thesis, University of Toronto, 2008.
- [17] Sertac Karaman and Emilio Frazzoli. Incremental sampling-based optimal motion planning. *submitted to Robotics: Science and Systems*, 2010.
- [18] L.E. Kavraki, P. Svestka, JC Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, August 1996.
- [19] J.J. Kuffner and S.M. LaValle. RRT-connect: An efficient approach to single-query path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 995–1001, 2000.
- [20] S. LaValle. Rapidly-exploring random trees: A new tool for path planning. Technical Report 98–11, Iowa State University, Dept. of Computer Science, 1998.
- [21] Steven M. LaValle, James J. Kuffner, and Jr. Randomized kinodynamic planning. *The International Journal of Robotics Research*, 20(5):378–400, 2001.
- [22] Colm ODunlaing. Motion planning with inertial constraints. *Algorithmica*, 2(4):431–475, 1987.
- [23] Gill A. Pratt. Low impedance walking robots. *Integ. and Comp. Biol.*, 42:174–181, 2002.
- [24] Gill A. Pratt, Matthew M. Williamson, Peter Dillworth, Jerry Pratt, Karsten Ulland, and Anne Wright. Stiffness isn't everything. In *Proceedings of the 4th International Symposium on Experimental Robotics (ISER)*, 1995.
- [25] Jerry Pratt. *Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots*. PhD thesis, Computer Science Department, Massachusetts Institute of Technology, 2000.
- [26] Zvi Shiller and Steven Dubowsky. Global time optimal motions of robotic manipulators in the presence of obstacles. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 370–375, 1988.
- [27] Alexander Shkolnik, Michael Levashev, Ian R. Manchester, and Russ Tedrake. Bounding on rough terrain with the littledog robot. *Under review*, 2010.
- [28] Ioan Alexandru Sucan. Kinodynamic motion planning for high-dimensional physical systems. Master's thesis, Rice University, Houston, TX, 2008.
- [29] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [30] Russ Tedrake. LQR-Trees: Feedback motion planning on sparse randomized trees. In *Proceedings of Robotics: Science and Systems (RSS)*, page 8, 2009.
- [31] Eduardo Torres-Jara. *Sensitive Manipulation*. PhD thesis, Massachusetts Institute of Technology, January 2007.
- [32] Konstantinos I. Tsianos, Ioan A. Sucan, and Lydia E. Kavraki. Sampling-based robot motion planning: Towards realistic applications. *Computer Science Review*, 1(1):2–11, August 2007.
- [33] Ronald van Ham, Thomas G. Sugar, Bram Vanderborgh, Kevin W. Hollander, and Dirk Lefeber. Compliant actuator designs: Review of actuators with passive adjustable compliance/controllable stiffness for robotic applications. *IEEE Robotics and Automation Magazine*, 16(3):81–94, September 2009.