

Collision-Free Humanoid Reaching: Past, Present, and Future

Evan Drumwright and Maja Matarić
University of Southern California
United States

Abstract

Most recent humanoid research has focused on balance and locomotion. This concentration is certainly important, but one of the great promises of humanoid robots is their potential for effective interaction with human environments through manipulation. Such interaction has received comparatively little attention, in part because of the difficulty of this task. One of the greatest obstacles to autonomous manipulation by humanoids is the lack of efficient collision-free methods for reaching. Though the problem of reaching and its relative, pick-and-place, have been discussed frequently in the manipulator robotics literature—e.g., (Lozano-Pérez et al., 1989); (Alami et al., 1989); (Burrige et al., 1995)—researchers in humanoid robotics have made few forays into these domains. Numerous subproblems must be successfully addressed to yield significant progress in humanoid reaching. In particular, there exist several open problems in the areas of algorithms, perception for modeling, and control and execution. This chapter discusses these problems, presents recent progress, and examines future prospects.

1 Introduction

Reaching is the one of the most important tasks for humanoid robots, endowing them with the ability to manipulate objects in their environment. Unfortunately, getting humanoids to reach efficiently and safely, without collision, is a complex problem that requires solving open subproblems in the areas of algorithms, perception for modeling, and control and execution. The algorithmic problem requires the synthesis of collision-free joint-space trajectories in the presence of moving obstacles. The perceptual problem, with respect to modeling, is comprised of acquiring sufficiently accurate information for constructing a geometric model of the environment. Problems of control and execution are concerned with correcting deviation from reference trajectories and dynamically modifying these trajectories during execution to avoid unexpected obstacles. This chapter delves into the relevant subproblems above in detail, describes the progress that has been made in solving them, and outlines the work remaining to be done in order to enable humanoids to perform safe reaching in dynamic environments.

2 Problem statement

The problem of reaching is formally cast as follows. Given:

1. a world $\mathcal{W} = \mathbb{R}^3$
2. the current time t_0 ; T is then defined as the interval $[t_0, \infty]$
3. a robot \mathcal{A}

4. a smooth manifold $\mathcal{X} \subset \mathbb{R}^n$ called the *state space* of \mathcal{A} ; let $\kappa : \mathcal{X} \rightarrow \mathcal{C}$ be a function that maps state-space to the robot's configuration space
5. the state transition equation \mathcal{A} is $\dot{x} = f(x, u, t)$, where $x \in \mathcal{X}$ and $u : T \rightarrow \mathcal{U}$ generates a vector of control inputs ($\mathcal{U} \subset \mathbb{R}^m$) as a function of time
6. a nonstationary obstacle region $\mathcal{O}(t) \subset \mathcal{W}, \forall t \in T; \mathcal{X}_{obs}(t), \forall t \in T$ is then the projection of obstacles in the robot's configuration space into state-space (i.e., $\kappa^{-1}(\mathcal{C}_{obs}(t)) = \mathcal{X}_{obs}(t)$) and $\mathcal{X}_{free}(t) = \mathcal{X} - \mathcal{X}_{obs}(t)$.
7. $\mathcal{D} \subset \mathbb{R}^6$ is the reachable workspace¹ of \mathcal{A}
8. a direct kinematics function, $F : \mathcal{C} \rightarrow \mathcal{D}$ that transforms robot states to operational-space configurations of one of the robot's end effectors
9. a set of feasible operational-space goal functions of time, G such that $\forall g \in G, g : T \rightarrow \mathcal{D}$
10. a feasible state-space Boolean function $\delta : T \times g \times \mathcal{X} \rightarrow \{0, 1\}$ where $g \in G$
11. $x_0 \in \mathcal{X}_{free}$, the state of the robot at t_0

generate the control vector function $u(\cdot)$ from time $t > t_0$ such that $x_t \in \mathcal{X}_{free}(t)$, where $x_t = x_0 + \int_{t_0}^t f(x, u(t), t) dt$, for $t > t_0$ and there exists a time t_j for which $\min_{g \in G} \|F(x(t_j)) - g(t_j)\| < \epsilon$ and $\delta(t_i, g(t_i), x(t_i)) = 1$ for all $t_i > t_j$, or correctly report that such a function $u(\cdot)$ does not exist.

Informally, the above states that to solve the reaching problem, the commands sent to the robot must cause it to remain collision-free and, at some point in the future, cause both the operational space distance from the end-effector to one of the goals to remain below a given threshold ϵ and the state-space of the robot to remain in an admissible region.

The implications of the above formal definition are:

- The state transition function $f(\cdot)$ should accurately reflect the dynamics of the robot. Unfortunately, due to limitations in mechanical modeling and the inherent uncertainty of how the environment might affect the robot, $f(\cdot)$ will only approximate the true dynamics. Section 4.3 discusses the ramifications of this approximation.
- The robot must have an accurate model of its environment. This assumption will only be true if the environment is instrumented or stationary. The environments in which humanoids are expected to operate are dynamic (see #6 above), and this chapter will assume that the environment is not instrumented. Constructing an accurate model of the environment will be discussed in Section 4.2.
- The goals toward which the robot is reaching may change over time; for example, the robot may refine its target as the robot moves nearer to it. Thus, even if the target itself is stationary, the goals may change given additional information. It is also possible that the target is moving (e.g., a part moving on an assembly line). The issue of changing targets will be addressed in Section 4.1.
- Manipulation is not explicitly considered. It is assumed that a separate process can grasp or release an object, given the operational-space target for the hand and the desired configuration for the fingers (the Boolean function $\delta(\cdot)$ is used to ensure that this latter condition is satisfied). This assumption is discussed further in the next section.

¹The *reachable workspace* is defined by Sciavicco & Siciliano (2000) to be the region of operational-space that the robot can reach with at least one orientation.

3 Related work

A considerable body of work relates to the problem defined in the previous section yet does not solve this problem. In some cases, researchers have investigated similar problems, such as developing models of human reaching. In other cases, researchers have attempted to address both reaching and manipulation. This section provides an overview of these alternate lines of research, though exhaustive surveys of these areas are outside of the scope of this chapter. Humanoids have yet to autonomously reach via locomotion to arbitrary objects in known, static environments, much less reach to objects without collision in dynamic environments. However, significant progress has been made toward solving these problems recently. This section concludes with a brief survey of methods that are directly applicable toward solving the reaching problem.

3.1 Models of reaching in neuroscience

A line of research in neuroscience has been devoted to developing models of human reaching; efficient, human-like reaching for humanoids has been one of the motivations for this research. Flash & Hogan (1985), Bullock et al. (1993), Flanagan et al. (1993), Crowe et al. (1998) and Thoroughman & Shadmehr (2000) represent a small sample of work in this domain. The majority of neuroscience research into reaching has ignored obstacle avoidance, so the applicability of this work toward safe humanoid reaching has not been established. Additionally, neuroscience often considers the problem of *pregrasping*, defined by Arbib et al. (1985) as a configuration of the fingers of a hand before grasping such that the position and orientation of the fingers with respect to the palm's coordinate system satisfies *a priori* knowledge of the object and task requirements. In contrast to the neuroscience approach, this chapter attempts to analyze the problem of humanoid reaching from existing subfields in robotics and computer science. Recent results in the domains of motion planning, robot mapping, and robot control architectures are used to identify remaining work in getting humanoids to reach safely and efficiently. This chapter is unconcerned with generating motion that is natural in appearance by using pregrasping and human models of reaching, for example.

3.2 Manipulation planning

Alami et al. (1997), Gupta et al. (1998), Mason (2001), and Okada et al. (2004) have considered the problem of *manipulation planning*, which entails planning the movement of a workpiece to a specified location in the world without stipulating how the manipulator is to accomplish the task. Manipulation planning requires reaching to be solved as a subproblem, even if the dependence is not explicitly stated. As noted in LaValle (2006), existing research in manipulation planning has focused on the geometric aspects of the task while greatly simplifying the issues of grasping, stability, friction, mechanics, and uncertainty. The reaching problem is unconcerned with grasping (and thereby friction) by presuming that reaching and grasping can be performed independently. The definition provided in Section 2 allows for treatment of mechanics (via $f(\cdot)$, the state transition function) and stability and uncertainty (by stating the solution to the problem in terms of the observed effects rather than the desired commands). Additionally, the problem of reaching encompasses more tasks than those used in pick-and-place; for example, both pointing and touching can be considered as instances of reaching.

A late development by Stilman & Kuffner, Jr. (2007) addresses manipulation planning amongst movable obstacles. The reaching problem as defined in Section 2 permits the obstacles to be movable by the humanoid. Many of the issues described in this chapter (e.g., constructing a model of the environment, monitoring execution, etc.) need to be resolved to fully explore this avenue, but the ability to move obstacles as necessary will certainly admit new solutions to some instances of the reaching problem.

3.3 Recent work directly applicable to humanoid reaching

Recent work in humanoid robotics and virtual humans is directly applicable toward efficient, safe humanoid reaching. For example, work by Brock (2000) permits reaching in dynamic environments by combining a planned path with an obstacle avoidance behavior. More recent work by Kagami et al. (2003) uses stereo vision to construct a geometric model of a static environment and motion planning and inverse kinematics to reach to and grasp a bottle with a stationary humanoid robot. Liu & Badler (2003); Kallmann et al. (2003); Bertram et al. (2006) and Drumwright & Ng-Thow-Hing (2006) focused on developing algorithms for humanoid reaching; the algorithms in the latter two works are *probabilistically complete*, while the algorithms in (Liu & Badler, 2003) and (Kallmann et al., 2003) are not complete in any sense. All four works assumed static environments, perfect control and holonomic constraints.

4 Outstanding issues

This section discusses the issues that remain to solve the reaching problem, as follows:

1. Constructing an accurate model of the environment
2. Planning collision-free motions in dynamic environments
3. Correcting deviation from the desired trajectory due to imperfect control
4. Avoiding both fixed and moving obstacles during trajectory execution

The first item has received the least research attention to date and therefore includes the majority of open problems in collision-free humanoid reaching. Section 4.2 discusses progress and prospects in this area. Planning collision-free motions, at least in static environments, has received considerable attention; Section 4.1 discusses why this problem is challenging from an algorithmic standpoint and addresses extensions to dynamic environments. Finally, correcting deviation from the planned trajectory and avoiding obstacles during trajectory execution are key to reach the target in a safe manner. Section 4.3 discusses these two issues.

4.1 Algorithmic issues

The best studied aspect of the reaching problem is the algorithmic component, which is an extension to the general motion planning introduced below. Section 4.1.1 formally relates the problem of reaching to the general motion planning problem, and analyzes the complexity of the latter. Section 4.1.2 introduces sample-based motion planning, a paradigm for circumventing the intractability of motion planning; the following section discusses the extension of a popular sample-based motion planner to respect differential constraints. Finally, Sections 4.1.4–4.1.8 discuss motion planning issues highly relevant to humanoid reaching, namely planning

under uncertainty, potential incompleteness resulting from multiple inverse kinematics solutions, planning to nonstationary goals, and planning in dynamic environments.

Researchers have investigated ways for planning collision-free motions from one configuration to another since the introduction of the *Piano Mover's Problem*, also known as the *Static Mover's Problem* Reif (1979). The Piano Mover's problem can be stated as follows (excerpted from LaValle (2006)).

Given:

1. a world \mathcal{W} , where $\mathcal{W} = \mathbb{R}^2$ or $\mathcal{W} = \mathbb{R}^3$
2. a semi-algebraic obstacle region $\mathcal{O} \subset \mathcal{W}$
3. a collection of m semi-algebraic links of a robot, $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m$
4. the configuration space \mathcal{C} of the robot; \mathcal{C}_{free} can then be defined as the subset of configuration space which does not cause the robot's geometry to intersect with any obstacles
5. the initial configuration of the robot, \mathbf{q}_I
6. the goal configuration of the robot, \mathbf{q}_G

generate a continuous path $\tau : [0, 1] \rightarrow \mathcal{C}_{free}$ such that $\tau(0) = \mathbf{q}_I$ and $\tau(1) = \mathbf{q}_G$ or correctly report that such a path does not exist. The term *semi-algebraic* refers to a geometric representation that is composed by Boolean operations on implicit functions.

4.1.1 Complexity of motion planning

The definition of the Piano Mover's Problem is quite similar to the problem formulation for reaching at the beginning of this chapter. Indeed, an instance of the reaching problem can be transformed into an instance Piano Mover's Problem given the following constraints:

- the obstacle region, \mathcal{O} is stationary (i.e., $\mathcal{O}(t_i) = \mathcal{O}(t_j), \forall t_i, t_j \in T$)
- $\mathcal{C} = \mathcal{X}$ from which follows $\kappa(x) \triangleq x, \forall x \in \mathcal{X}$ and $\mathbf{x}_0 = \mathbf{q}_I$
- $\mathcal{U} = \mathcal{X}$
- G consists of a single element, g , which is nonstationary, and there exists only one robot configuration \mathbf{q}_G that results in g (i.e., $F^{-1}(g) = \mathbf{q}_G$)
- $\int_{t_a}^{t_b} f(x, u(t), t) dt = u(t_b) - x(t_a)$ (implies that the command is the new robot state)

Reif (1979) showed that the Piano Mover's Problem is PSPACE-complete, implying NP-hardness. Additionally, the best known algorithm for solving the Piano Mover's problem (complexity-wise) is Canny's Roadmap Algorithm (Canny, 1993), which exhibits running-time exponential in the configuration space; aside from being intractable, the algorithm is reportedly quite difficult to implement (LaValle, 2006). Later work by Reif & Sharir (1994) proved that planning motions for a robot with fixed degrees-of-freedom and velocity constraints in the presence of moving obstacles with known trajectories is PSPACE-hard; thus, the constraints that were imposed transforming the reaching problem into the Piano Mover's Problem are unlikely to make the former problem easier computationally.

4.1.2 Sample-based motion planning

The paradigm of *sample-based motion planning* was introduced with the advent of the *randomized potential field* (Barraquand & Latombe, 1991). Sample-based algorithms trade algorithmic completeness for excellent average-case performance and ease of implementation. In fact, completeness was not cast aside; rather, it was relaxed to lesser constraints, *probabilistic completeness* and *resolution completeness*. It is said that an algorithm is probabilistically complete if the probability of finding a solution, if it exists, tends to unity as the number of samples increase. Similarly, a motion planning algorithm is resolution complete if a solution, if it exists, will be found in finite time given sufficiently dense sampling resolution over the domain of configuration space. Note that neither weaker form of completeness permits this class of algorithm to definitively state that a path does not exist. Finally, the underlying approach of sample-based planning is quite different from adapting classical search algorithms (i.e., A^* , best-first, etc.) to a discretized grid; such an approach is generally intractable due to the combinatorial explosion of configuration space.

The most popular sample-based algorithm is currently the rapidly-exploring random tree (RRT) LaValle (1998). RRTs are popular due to their inherent bias toward the largest Voronoi region of configuration space (i.e., the largest unexplored region) during exploration. Efficient exploration is critical for sample-based algorithms because their running times generally increase linearly with the number of samples.

4.1.3 Motion planning under differential constraints

In addition to the advantage of efficient exploration, RRTs allow for planning under differential (e.g., nonholonomic) constraints, through *kinodynamic planning*. Kinodynamic planning plans in the control space, rather than in configuration space, and is therefore able to respect dynamic constraints. Kinodynamic planning theoretically permits motion to be planned for humanoids, which generally use bipedal locomotion and are nonholonomically constrained. As might be expected, kinodynamic planning is harder computationally than in the unconstrained case; additionally, kinodynamic planning requires a model of the system's dynamics and a control method for solving a *two-point boundary value problem*.²

Planning directly in the state-actuator space of the humanoid is infeasible: the motion planner would not only have to avoid obstacles but also balance the humanoid and provide locomotion. An accurate model of the robot's dynamics would be required as well. Alternatively, planning could treat the robot's base as a three degree-of-freedom virtual joint corresponding to the manifold $SE(2)$, on which solving the two-point boundary value problem is easy. This was the approach taken in Drumwright (2007). Once a plan is formulated in this new space, a mechanism is required to determine kinematic commands for the robot's locomotory degrees-of-freedom (e.g., legs, wheels, etc.) that yield the kinematic commands for the virtual base.

Finally, it should be noted that if humanoids moved on holonomic bases, simpler methods could potentially be employed for humanoid reaching. For example, Maciejewski & Klein (1985) combines inverse kinematics with obstacle avoidance for redundant manipulators under holonomic constraints; some mobile manipulators fit into this category. Though the method of

²In the context of this work, the two-point boundary problem can be considered to be the problem of getting from a given state to a desired state under nonholonomic constraints. For example, the problem of parallel parking a car can be considered a two-point boundary problem: a series of actions is required to move the car into a parking space, even if only a simple translation is required (e.g., the car initially is aligned laterally with the parking space).

Maciejewski and Klein can cause the robot to become trapped in local minima, its computational requirements are minimal.

4.1.4 Motion planning for humanoid reaching

The RRT has been applied successfully to reaching for humanoids in virtual environments in (Kuffner, Jr., 1998); (Liu & Badler, 2003); (Kallmann et al., 2003); (Bertram et al., 2006) and (Drumwright & Ng-Thow-Hing, 2006), among others. Additionally, the RRT has been applied to reaching for an embodied humanoid by Kagami et al. (2003), although, as stated in Section 2, the environment was considered to be stationary and locomotion was not utilized.

Unfortunately, several assumptions of current, effective motion planning algorithms are unsuitable for humanoid reaching, as follows:

1. An accurate model of the environment and a precise mechanism for control are required.
2. A single goal configuration is assumed.
3. Targets for reaching do not change over time.
4. The environment is static.

These issues are explored further in the remainder of this section.

4.1.5 Planning under uncertainty

This chapter generally assumes that a sufficiently accurate model of the environment can be constructed and that the robot can be controlled with some degree of accuracy. However, these assumptions are often unrealistic in real-world robotics. It would be advantageous to construct plans that would maximize the distance from the robot to obstacles, for example, to minimize deleterious effects of uncertainty. LaValle & Hutchinson (1998) have explored the issue of planning under sensing and control uncertainty; however, their use of dynamic programming has restricted applications to planar robots due to the “curse of dimensionality” (Bellman, 1957). As an alternative, Lazanas & Latombe (1995) proposed an approach based on *landmarks*, regions of the state space where sensing and control are accurate. The assumption that such regions exist is significant. The limited application of these two methods illustrates the difficulty of planning while maximizing objectives, which is known as *optimal planning* (LaValle, 2006). Finally, Baumann (2001) proposes a planning method that iteratively modifies a trajectory; a fitness function judges the quality of the modified trajectory versus the original, in part based on distance to obstacles. However, this work has yet to be subjected to peer-review.

4.1.6 Incompleteness resulting from multiple IK solutions

A single operational-space goal generally corresponds to an infinite number of robot configurations, given the hyper-redundant degrees-of-freedom of most humanoids.³ It is possible that collision-free paths exist only for a subset of the space of collision-free inverse kinematics solutions. Drumwright & Ng-Thow-Hing (2006) addressed that problem by maintaining a list of goal configurations that is continually grown using inverse kinematics; the motion planner frequently attempts to reach arbitrary goals in the list. The motion planner avoids becoming stuck planning to unreachable goals by not committing to any particular goal.

³We assume that all available degrees-of-freedom are used for reaching, rather than for achieving secondary tasks like singularity avoidance, e.g., (Tanner & Kyriakopoulos, 2000).

4.1.7 Motion planning to nonstationary goals

Planning trajectories to nonstationary goals has received little attention in the motion planning community; however, two concerns are evident. The goals may change with increasing sensory data, leading to new estimates of the target's position and orientation. Additionally, the target itself may be moving, perhaps with a predictable trajectory.

The first issue is readily solvable using existing methods. If already in the process of planning, the goal can be replaced without ill effects: the results from sample-based planning methods will not be invalidated. If a plan has been determined, replanning can be performed using the results already determined (i.e., the roadmap, tree, etc.) to speed computation. Alternatively, inverse kinematics can be utilized to modify the generated plan slightly, somewhat in the manner used for motion retargeting by Choi & Ko (2000).

Moving targets are far more difficult to manage. Current sample-based planning methods have not focused on this problem, and the complexity of adapting existing methods to this purpose is unknown. Again, it is imaginable that existing plans could be modified using inverse kinematics, though replanning may be required if the target is moving quickly. Alternatively, Brock's method (discussed in Section 4.3) could potentially be utilized with some modifications toward this purpose.

4.1.8 Motion planning in dynamic environments

A simple means to address motion planning in dynamic environments adds time as an extra dimension of configuration-space. As LaValle (2006) notes, the planning process must constrain this extra dimension to move forward only. This approach of augmenting the configuration space with time can fail because the dynamics of the robot are not considered: it may not be possible or advisable for a robot to generate sufficient forces to follow the determined trajectory. An alternative to this approach is to use kinodynamic planning, as described in Section 4.1.3 to plan in the control space of the robot. However, kinodynamic planning does not appear to be applicable to configuration spaces above twelve dimensions, in addition to the difficulties with this approach described in Section 4.1.3.

Dynamic replanning, which refers to fast replanning as needed, is an alternative to methods which plan around dynamic obstacles. Dynamic replanning can treat dynamic obstacles as stationary and thus avoids the additional complexity of planning around these obstacles. Dynamic replanning may be the best option for the high-dimensional configuration spaces of humanoids. Kallmann & Mataric (2004) have explored online modification of existing sample-based roadmaps for virtual humanoids; unfortunately, that work indicated that the modification is likely no faster than generating a new plan. However, newer algorithms by Ferguson & Stentz (2006) and Zucker et al. (2007), also based on RRTs, have proven adept at replanning in nonstationary environments with large configuration spaces by modifying existing trees dynamically. Additionally, Drumwright & Ng-Thow-Hing (2006) have indicated that even planning anew in real-time for humanoids using a slightly modified RRT algorithm is nearly performable using current computers.

Recent work by Jaillet & Siméon (2004) and van den Berg & Overmars (2005) has taken a two-phase approach to motion planning in dynamic environments. The first phase entails constructing a probabilistic roadmap (Kavraki et al., 1996) over the persistent parts of the environment offline. In the online phase, a graph search algorithm finds a feasible path around dynamic obstacles. These methods are compelling because the bulk of computation, construct-

ing the roadmap, is performed offline. Nevertheless, further research is required to determine efficient ways to update the roadmap as the environment changes (e.g., doors are opened, furniture is moved, etc.) before these methods can be used for humanoid reaching,

Regardless of the method employed, a significant concern is that the process of planning can itself cause possible solutions to disappear, because planning occurs in real-time. An additional significant concern is the requirement of many methods that the trajectories of the obstacles to be known *a priori*; filtering techniques—e.g., Kalman filters Kalman & Bucy (1961)—might permit short-term predictions, but long-term predictions will be problematic unless the obstacles follow simple trajectories.

4.1.9 Summary

Substantial progress has been made in motion planning in the past decade, leading to tractable solutions of high-dimensional planning problems. Indeed, in the areas of planning to nonstationary goals and planning in dynamic environments, researchers are on the cusp of solutions that are viable for humanoid reaching. However, considerable work still remains in the area of planning under uncertainty.

4.2 Perception for modeling issues

We ignore the problems of object recognition and object pose determination, which are beyond the scope of this chapter, and focus on the perceptual issues related to modeling the environment for purposes of collision detection⁴, assuming that the humanoid is equipped with an accurate range finding sensor. The problem of simultaneous localization and mapping (SLAM) is well studied in the mobile robotics community, where significant success has been achieved at developing methods to construct maps of $2\frac{1}{2}$ D environments; some success has been achieved building three-dimensional maps as well. The natural inclination is to see whether SLAM techniques for mobile robots can be adapted to the problem of environment modeling toward humanoid reaching. Human environments are dynamic, humanoids manipulate objects (thus changing the world), and environment modeling is conducted with respect to planning; these challenges make current SLAM methods for mobile robotics difficult to utilize for humanoid reaching. Indeed, significant obstacles remain before humanoids can construct environment models suitable for motion planning. The remainder of this section discusses the relevant issues toward this purpose, namely representation, localization, exploration, and nonstationary environments.

4.2.1 Representation of environment model

There exist several possible representations for modeling the environment, including volumetric point sets (Thrun et al., 2003); occupancy grids (Moravec & Elfes, 1985), (Elfes, 1989); 3D models (Teller, 1998), (Kagami et al., 2003) and feature-based maps (Kuipers et al., 1993). However, the application of reaching presents several requirements. First, the representation must allow for fast intersection testing. Second, the representation must be able to efficiently manage the prodigious amounts of data that 3D range scans generate. These first two requirements

⁴Although the goal functions for the reaching problem are given in the global coordinate frame, it is quite natural to use ego-centric frames instead. As a result, localization is likely not required to perform reaching, except for effectively constructing the environment model.

exclude volumetric point sets. Fast updates of the representation from sensory data are also required. This stipulation excludes the use of 3D models, which require considerable post-processing including iso-surface extraction or stitching (Turk & Levoy, 1994) and aligning and registering (Mayer & Bajcsy, 1993); (Pito, 1996). In addition to the requirements listed above, the ability to readily extract high-level features from the representation for use with localization, described in Section 4.2.2, would be advantageous.

Drumwright et al. (2006) used an octree, which may be considered as an extension of occupancy grids, for modeling environments for reaching. The octree representation results in efficient storage, permits fast updates to the representation from range data, and is capable of performing fast intersection testing with the robot model. Drumwright et al. (2006) assumed perfect localization, and we are unaware of work that extracts features from octrees (or their two-dimensional equivalent, quadtrees) for the purpose of localization. However, there is precedent for feature extraction from octrees, e.g., (Sung et al., 2002).

One possible alternative to the octree representation would use high-level features (e.g., landmarks, objects, or shapes) as the base representation. Such features would serve well for input to one of the popular methods for simultaneous localization and mapping (SLAM) such as (Smith & Cheeseman, 1985); (Smith et al., 1990) or (Fox et al., 1999). Recognition of objects in the context of mapping and localization has been limited generally to those objects which are of interest to mobile robots, including doors (Avots et al., 2002), furniture (Hähnel et al., 2003), and walls (Martin & Thrun, 2002). Additionally, representations used typically involve geometric primitives, which fail to realize the potential benefit of using identified objects to infer occluded features. The difficulty of performing object recognition in unstructured environments makes the near-term realization of this benefit unlikely.

4.2.2 Localization

The humanoid must be able to localize itself (i.e., know its planar position and orientation) with respect to the environment, if not globally. Recent work by Fox et al. (1999) indicates that localization can be performed successfully with range data even in highly dynamic environments. This work, as well as other probabilistic approaches, typically can provide measures of certainty of their localization estimates. The estimated variance can be used to filter updates of the environment model; the environment model will be updated only if the certainty of the prediction is above a given threshold. Alternatively, localization accuracy can be quite high, e.g., on the order of centimeters (Yamamoto et al., 2005), if the environment is modestly instrumented.

4.2.3 Nonstationary environments

Modeling nonstationary environments requires the management of three types of obstacles, which we call *dynamic*, *movable*, and *static*. Dynamic obstacles are capable of moving on their own (e.g., humans, cars, etc.), while movable obstacles (e.g., cups, furniture, books, etc.) must be moved by a dynamic obstacle. Static obstacles, such as walls, are incapable of movement. The three cases could be managed separately or movable and static obstacles could be treated identically yet permit gradual changes to the environment.

Three recent approaches have made strides toward modeling nonstationary environments. The first approach, introduced independently by Wang & Thorpe (2002) and Hähnel et al. (2002), attempts to identify dynamic objects and filter them from the sensory data. The second approach, (Hähnel et al., 2003), uses an expectation-maximization (Dempster et al., 1977)

based algorithm to repeatedly process the sensor readings and predict whether a given feature belongs to a dynamic obstacle; dynamic obstacles are not incorporated into the constructed map. Finally, Biswas et al. (2002) attempt to model dynamic obstacles as movable rigid bodies using an occupancy grid and an expectation-maximization based algorithm. Unfortunately, the first two approaches described above are ill-suited for movable obstacles, while the second approach is unable to handle dynamic obstacles.

4.2.4 Exploration

Exploration of the environment to facilitate reaching requires both directed locomotion and directed/active sensing. Both exploration and active sensing have been studied extensively in the context of mobile robot mapping, it is unlikely that this research is fully applicable toward our problem. Exploration in the mapping context is conducted in a greedy manner: movement is directed to build a complete map of the environment with respect to some criterion, e.g., minimum time (Yamauchi, 1997), minimum energy expenditure (Mei et al., 2006), maximum safety (González-Baños & Latombe, 2002), etc. In contrast, the reaching problem requires a balance between exploration and exploitation: the environment must be sufficiently explored, but not at the expense of finding a valid collision-free path. The balance between exploration and exploitation precludes the use of active sensing frameworks, such as that described by Mihaylova et al. (2002), to guide exploration. Finally, moving obstacles must be avoided during exploration, perhaps by using reactive approaches like VFH Borenstein & Koren (1989) or potential fields Khatib (1986).

Choset & Burdick (2000) refer to environmental exploration in the context of motion planning *sensor based motion planning*, and introduced a new data structure, the *hierarchical generalized voronoi graph*, to solve this problem. Unfortunately, their work was targeted toward mobile robots, and it appears highly unlikely to scale to humanoid robots with high-dimensional configuration spaces.

Researchers have yet to propose methods to perform directed exploration with respect to motion planning for humanoid robots. One possibility is to adapt the concepts employed for mobile robot exploration to humanoids. For example, *frontier-based exploration* (Yamauchi, 1997), which directs the robot to move to the area between open space and uncharted territory (i.e., the frontier), could be combined with a heuristic to select frontiers nearest the target.

4.2.5 Summary

This section covered the relevant issues necessary to construct a model of the environment for reaching. The issues of localization and selecting a proper representation for modeling the environment seem manageable through existing research. However, considerable work remains in the areas of modeling nonstationary environments and exploration for motion planning.

4.3 Control and Execution issues

Ideally, kinematic commands could be issued to a humanoid, the robot would execute those commands with perfect precision, and obstacles would not move. Those assumptions do not hold in the real-world: humanoids are typically holonomically constrained, controllers are imperfect, robots exhibit some degree of inaccuracy due to mechanical tolerances, the model of the environment may be flawed, and obstacles can appear suddenly.

The process of getting a manipulator to reach to an operational space target generally consists of the following sequence of steps:

1. A motion planner generates kinematic commands; the environment is often considered to be static for manipulators, so motion planning may consist of using cubic or quintic splines to form a trajectory with via points.
2. The kinematic commands are input to a mechanism for rescaling the trajectory. The trajectory is rescaled to a longer duration, if necessary, to permit the robot to execute the trajectory without deviating from the prescribed path and prevent possible damage to the robot.
3. The commands are continually fed to a controller, which translates kinematic commands into actuator commands.

This process runs counter to best practices in mobile robotics because it provides no means to monitor the execution of the trajectory and react to exigent events, such as humans getting in the way. For humanoids, the process above must provide a means to supervise the robot's execution of the plan to avoid obstacles and correct deviation from the plan.

This section discusses the causes, effects, and remedies for controller deviation and examines execution monitoring, which attempts to avoid obstacles while driving the robot to its target.

4.3.1 Controller deviation

Robot controllers that minimize command error generally incorporate a model of the robot's dynamics; in particular, an inverse dynamics model of the robot is utilized to map desired accelerations to actuator torques. Controllers that use inverse dynamics models will still exhibit some deviation from the desired trajectory due to approximations of the robot's true dynamics (e.g., Coulombic friction, infinitely rigid links, etc.) Of far more concern, however, is that an algorithm does not exist for computing the inverse dynamics for robots with floating-bases experiencing contact constraints (e.g., for humanoids standing on a floor). As a result, humanoids are frequently controlled using feedback controllers (e.g., PD, PID, PIDD², etc.) These controllers necessarily result in deviation from the commanded trajectory because their generated motor commands are solely a function of command error.

For systems with holonomic differential constraints, deviation from the reference trajectory is the primary concern. For systems that are nonholonomically constrained, such as bipeds, deviation results in an additional difficulty: correcting deviation may require solving a two-point boundary value problem. Thus, replanning may be necessary to correct for deviation of the robot's base.

4.3.2 Global methods for execution monitoring

The execution of the planned trajectory must be monitored to avoid both dynamic obstacles and static obstacles that were circumlocuted during planning but, due to unintentional deviation from the trajectory, that the robot is in imminent danger of contacting. Ideally, a policy would exist that could issue avoidance commands to the robot with expediency. Such a policy, denoted as $\mathbf{u} \leftarrow \pi(\mathbf{q}, \dot{\mathbf{q}}, \mathcal{O})$, would issue commands (i.e., \mathbf{u}) as a function of the current state of the robot (q, \dot{q}) and the current obstacle region, \mathcal{O} . This policy theoretically could be constructed offline through reinforcement learning or dynamic programming; practically, however,

the high-dimensionality of the robot's state-action spaces and the infinite number of configurations of the environment make this approach intractable (Bellman, 1957). Alternatively, it is conceivable that a policy could be determined nearly instantaneously using reinforcement learning as the environmental model changes (i.e., $\pi(\cdot)$ would be a function of the robot's configuration only). Again, such an approach is currently infeasible given the robot's high-dimensional state-action spaces, and dynamic replanning would likely prove to be a better alternative. Yang & LaValle (2000) provide a method for *feedback motion planning*, which entails constructing motion plans that are viable from entire regions of state-space. Perhaps because objective criteria, upon which dynamic programming is based, do not have to be considered, their method seems to be slightly more tractable than dynamic programming. However, this method has only been applied to five-dimensional configuration spaces, and it appears unlikely to scale much further.

The methods described above are known as "global" methods, because good actions (i.e., those that lead to the goal and away from obstacles) are always generated, regardless of the configurations of the robot and environment. As noted above, the combined state-action spaces are immense, thus generally preventing global methods from being used for humanoids. "Local methods", in contrast to global methods, are generally tractable, but may result in failure to reach the goal; fortunately, these methods typically work well at avoiding obstacles.

4.3.3 Local methods for execution monitoring

One local method, by Brock & Khatib (2002), currently permits reaching with fast obstacle avoidance. This method is based on *elastic strips*, an extension of Khatib's operational space framework (Khatib, 1986) that employs virtual rubber bands. These virtual rubber bands warp planned trajectories away from obstacles. The most significant drawback to using elastic strips arises from its use of the robot's joint-space inertia matrix; this matrix is frequently unknown—it requires precise knowledge of the robot's inertial properties—and changes if the robot carries objects or interacts with the environment. In addition, like all local methods, elastic strips can cause the robot to become stuck in local minima. However, Brock did implement this method on both a mobile manipulator in an environment populated with moving obstacles and a simulated humanoid.

4.3.4 Summary

The ability of various humanoids to accurately locomote along prescribed paths in Cartesian space is currently unknown. However, control and locomotion are both active research domains in humanoid robotics. Continuing research might render deviation from planned trajectories an insignificant issue for humanoid reaching in the future, particularly if robots can be controlled with high accuracy in a black-box manner.

For execution monitoring, only elastic strips has proven successful for real-time obstacle avoidance on a mobile manipulator that followed a given trajectory. Depending on how seriously inaccuracies in the joint-space inertia matrix affect elastic strips, this method might prove extendible to humanoids. Viable alternatives include dynamic replanning and the offline-online probabilistic roadmap based methods of Jaillet & Siméon (2004) and van den Berg & Overmars (2005) described in Section 4.1, using dense roadmaps of configuration space. Regardless of the method chosen for execution monitoring, dynamic obstacles still require representation in the robot's perceptual model of the environment, as described in Section 4.2.

5 Conclusion

Table 1: The core issues with respect to humanoid reaching discussed in this chapter with brief summaries of the findings.

Issue	Progress and prospects
Planning under uncertainty	Difficult for humanoids, possible solution is dynamic replanning
Incompleteness resulting from multiple IK solutions	Solved by (Bertram et al., 2006) and (Drumwright & Ng-Thow-Hing, 2006)
Planning to nonstationary goals	Has received little focus, possible solution is modification of existing plans using inverse kinematics
Planning in dynamic environments	Dynamic replanning and offline-online probabilistic roadmaps seem to be promising directions; effectiveness toward humanoid reaching needs to be assessed
Representation of environment model	Octree approach works well, high-level features might work better in the future
Localization	Highly active area of research; (Fox et al., 1999) indicates good localization feasible even in nonstationary environments
Modeling nonstationary environments	Considerable work remaining; current work fails to address movable obstacles and dynamic obstacles simultaneously
Exploration	Considerable work remaining; exploration for planning humanoid motions yet to be addressed
Controller deviation	Better controllers will address this problem
Execution monitoring	(Brock & Khatib, 2002) is one possible solution; dynamic replanning is another

The past decade has seen remarkable progress in the types of motion planning problems that have been solved. However, many of these successes have been in application domains far removed from robotics. Additionally, humanoids have yet to perform tasks autonomously in human environments. This chapter formally defined one of the most important problems for humanoids, efficient, safe reaching in dynamic, unknown environments. Progress in the three areas critical areas to solving this problem—algorithms, perception for modeling, and execution—was surveyed and future prospects were examined. Table 1 summarizes the key issues in these areas as well as progress and prospects.

With respect to algorithms, control, and execution monitoring, dynamic replanning seems to be one of the best solutions for humanoid reaching. As computing power continues to grow, dynamic replanning becomes increasingly viable for planning reaching in nonstationary environments. For modeling the robot’s environment, computational power seems to be less of an

issue. Instead, new algorithms are required to balance the tradeoff between exploration and exploitation and to perceive and classify fully dynamic and semi dynamic obstacles. These three areas critical to humanoid reaching are currently the focus of intensive research, and results applicable toward humanoid reaching are on the horizon.

References

- Alami, R., Laumond, J.-P., & Siméon, T. (1997). Two manipulation planning algorithms. In J.-P. Laumond & M. Overmars (Eds.), *Algorithms for robot motion and manipulation*. Wellesley, MA: A.K. Peters.
- Alami, R., Siméon, T., & Laumond, J.-P. (1989). A geometric approach to planning manipulation tasks. In *Proc. intl. symp. on robotics research (ISRR)* (pp. 113–119).
- Arbib, M. A., Iberall, T., & Lyons, D. (1985). Coordinated control programs for movements of the hand. *Experimental Brain Research*, 111–129.
- Avots, D., Lim, E., Thibaux, R., & Thrun, S. (2002). A probabilistic technique for simultaneous localization and door state estimation with mobile robots in dynamic environments. In *Proc. of the IEEE conf. on intelligent robots and systems (IROS)*. Lausanne, Switzerland.
- Barraquand, J., & Latombe, J.-C. (1991, Dec). Robot motion planning: a distributed representation approach. *Intl. J. of Robotics Research*, 10(6), 628–649.
- Baumann, A. (2001). *Robot motion planning in time-varying environments*. Unpublished doctoral dissertation.
- Bellman, R. E. (1957). *Dynamic programming*. Dover Publications.
- Bertram, D., Kuffner, Jr., J. J., Dillmann, R., & Asfour, T. (2006, May). An integrated approach to inverse kinematics and path planning for redundant manipulators. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)*. Orlando, FL, USA.
- Biswas, R., Limketkai, B., Sanner, S., & Thrun, S. (2002). Towards object mapping in dynamic environments with mobile robots. In *Proc. of the IEEE conf. on intelligent robots and systems (IROS)*. Lausanne, Switzerland.
- Borenstein, J., & Koren, Y. (1989, Sept–Oct). Real-time obstacle avoidance for fast mobile robots. *IEEE Trans. on Systems, Man, and Cybernetics*, 19(5), 1179–1187.
- Brock, O. (2000). *Generating robot motion: the integration of planning and execution*. Unpublished doctoral dissertation, Stanford University.
- Brock, O., & Khatib, O. (2002, Dec). Elastic strips: A framework for motion generation in human environments. *Intl. J. of Robotics Research*, 21(12), 1031–1052.
- Bullock, D., Grossberg, S., & Guenther, F. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *J. of Cognitive Neuroscience*, 5, 408–435.
- Burrige, R. R., Rizzi, A. A., & Koditschek, D. E. (1995, Aug). Toward a dynamical pick and place. In *Proc. of the 1995 IEEE/RSJ intl. conf. on intelligent robots and systems* (Vol. 2, pp. 292–297).

- Canny, J. (1993). Computing roadmaps of general semi-algebraic sets. *The Computer Journal*, 36(5), 504–514.
- Choi, K.-J., & Ko, H.-S. (2000). On-line motion retargeting. *Journal of Visualization and Computer Animation*, 11, 223–235.
- Choset, H., & Burdick, J. (2000). Sensor based motion planning: The hierarchical generalized voronoi graph. *Intl. J. of Robotics Research*, 19(2), 96–125.
- Crowe, A., Porrill, J., & Prescott, T. (1998, Sept). Kinematic coordination of reach and balance. *Journal of Motor Behavior*, 30(3), 217-233.
- Dempster, A. P., Laird, A. N., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. of the Royal Statistical Society, Series B*, 39(1), 1–38.
- Drumwright, E. (2007). *The task matrix: a robot-independent framework for programming humanoids*. Unpublished doctoral dissertation.
- Drumwright, E., & Ng-Thow-Hing, V. (2006, Oct). Toward interactive reaching in static environments for humanoid robots. In *Proc. of IEEE/RSJ intl. conf. on intelligent robots and systems (IROS)*. Beijing.
- Drumwright, E., Ng-Thow-Hing, V., & Matarić, M. (2006, Dec). The task matrix framework for platform independent humanoid programming. In *Proc. of IEEE intl. conf. on humanoid robotics (humanoids)*. Genova, Italy.
- Elfes, A. (1989, June). Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6), 46–57.
- Ferguson, D., & Stentz, A. (2006, May). Replanning with RRTs. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)*. Orlando, FL.
- Flanagan, J. R., Feldman, A. G., & Ostry, D. J. (1993). Control of trajectory modifications in target-directed reaching. *Journal of Motor Behavior*, 25(3), 140-152.
- Flash, T., & Hogan, N. (1985, July). The coordination of arm movements: An experimentally confirmed mathematical model. *The Journal of Neuroscience*, 5(7), 1688-1703.
- Fox, D., Burgard, W., & Thrun, S. (1999). Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, 11, 391–427.
- González-Baños, H., & Latombe, J.-C. (2002, Oct–Nov). Navigation strategies for exploring indoor environments. *Intl. J. of Robotics Research*, 21(10–11), 829–848.
- Gupta, K., Ahuactzin, J., & Mazer, E. (1998). Manipulation planning for redundant robots: A practical approach. *Intl. J. of Robotics Research*, 17(7).

- Hähnel, D., Schulz, D., & Burgard, W. (2002, Sept–Oct). Mapping with mobile robots in populated environments. In *Proc. of the IEEE/RSJ intl. conf. on intelligent robots and systems (IROS)*. Lausanne, Switzerland.
- Hähnel, D., Triebel, R., Burgard, W., & Thrun, S. (2003, Sept). Map building with mobile robots in dynamic environments. In *Proc. of the IEEE intl. conf. on robotics and automation*. Taipei, Taiwan.
- Jaillet, L., & Siméon, T. (2004, Sept–Oct). A PRM-based motion planner for dynamically changing environments. In *Proc. of the IEEE conf. on intelligent robots and systems (IROS)* (Vol. 2, pp. 1606–1611). Sendai, Japan.
- Kagami, S., Kuffner, Jr., J. J., Nishiwaki, K., Inaba, M., & Inoue, H. (2003, April). Humanoid arm motion planning using stereo vision and RRT search. *Journal of Robotics and Mechatronics*, 15(2), 200-207.
- Kallmann, M., Aubel, A., Abaci, T., & Thalmann, D. (2003). Planning collision-free reaching motions for interactive object manipulation and grasping. In *Proc. of eurographics* (p. 313-322). Grenada, Spain.
- Kallmann, M., & Matarić, M. (2004, April). Motion planning using dynamic roadmaps. In *Proc. of IEEE intl. conf. on robotics and automation (ICRA)* (p. 4399-4404). New Orleans, LA.
- Kalman, R., & Bucy, R. (1961). New results in linear filtering and prediction theory. *Trans. of the ASME – Journal of Basic Engineering*, 83, 95–107.
- Kavraki, L. E., Svestka, P., Latombe, J.-C., & Overmars, M. (1996). Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Trans. on Robotics and Automation*, 12(4), 566-580.
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The Intl. J. of Robotics Research*, 5(1), 90–98.
- Kuffner, Jr., J. J. (1998). Goal-directed navigation for animated characters using real-time path planning and control. *Lecture Notes in Computer Science*, 1537, 171-186.
- Kuipers, B., Froom, R., Lee, W. K., & Pierce, D. (1993). The semantic hierarchy in robot learning. In J. Connell & S. Mahadevan (Eds.), *Robot learning*. Boston, MA: Kluwer Academic Publishers.
- LaValle, S. (1998, October). *Rapidly-exploring random trees: a new tool for path planning* (Tech. Rep. No. TR 98-11). Computer Science Dept., Iowa State University.
- LaValle, S. (2006). *Planning algorithms*. Cambridge University Press.
- LaValle, S., & Hutchinson, S. (1998, Jan). An objective-based framework for motion planning under sensing and control uncertainties. *Intl. J. of Robotics Research*, 17(1), 19–42.

- Lazanas, A., & Latombe, J. C. (1995). Motion planning with uncertainty: A landmark approach. *Artificial Intelligence*, 76(1-2), 285-317.
- Liu, Y., & Badler, N. I. (2003, May). Real-time reach planning for animated characters using hardware acceleration. In *Proceedings of computer animation and social agents* (p. 86-93). New Brunswick, NJ: IEEE Computer Society.
- Lozano-Pérez, T., Jones, J. L., Mazer, E., & O'Donnell, P. A. (1989). Task-level planning of pick-and-place robot motions. *Computer*, 22(3), 21-29.
- Maciejewski, A. A., & Klein, C. A. (1985). Obstacle avoidance for kinematically redundant manipulators in dynamically varying environments. *Intl. J. of Robotics Research*, 4(3), 109-117.
- Martin, C., & Thrun, S. (2002). Real-time acquisition of compact volumetric maps with mobile robots. In *IEEE intl. conf. on robotics and automation (ICRA)*. Washington, DC.
- Mason, M. T. (2001). *Mechanics of robotic manipulation*. Cambridge, MA: MIT Press.
- Mayer, J., & Bajcsy, R. (1993, May). Occlusions as a guide for planning the next view. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(5), 417-433.
- Mei, Y., Lu, Y.-H., Hu, Y. C., & Lee, C. S. G. (2006). Energy-efficient mobile robot exploration. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)*. Orlando, FL.
- Mihaylova, L., Lefebvre, T., Bruyninckx, H., Gadeyne, K., & De Schutter, J. (2002). Active sensing for robotics—a survey. In *Intl. conf. on numerical methods and applications*.
- Moravec, H., & Elfes, A. (1985, March). High-resolution maps from wide-angle sonar. In *Proc. of IEEE intl. conf. on robotics and automation (ICRA)*. St. Louis, MO.
- Okada, K., Haneda, A., Nakai, H., Inaba, M., & Inoue, H. (2004, Sept-Oct). Environment manipulation planner for humanoid robots using task graph that generates action sequence. In *Proc. of the IEEE/RSJ intl. conf. on intelligent robots and systems (IROS)*. Sendai, Japan.
- Pito, R. (1996). A sensor based solution to the next best view problem. In 941-945 (Ed.), *Proc. IEEE intl. conf. on pattern recognition* (Vol. 1).
- Reif, J. H. (1979). Complexity of the mover's problem and generalizations. In *IEEE symposium on foundations of computer science* (pp. 421-427). San Juan, Puerto Rico.
- Reif, J. H., & Sharir, M. (1994). Motion planning in the presence of moving obstacles. *Journal of the Association of Computing Machinery*, 41, 764-790.
- Sciavicco, L., & Siciliano, B. (2000). *Modeling and control of robot manipulators, 2nd ed.* London: Springer-Verlag.

- Smith, R., & Cheeseman, P. (1985). *On the representation and estimation of spatial uncertainty* (Tech. Rep. Nos. TR 4760, TR 7239). SRI.
- Smith, R., Self, M., & Cheeseman, P. (1990). Estimating uncertain spatial relationships in robotics. In I. Cox & G. Wilfong (Eds.), *Autonomous robot vehicles* (pp. 167–193). Springer-Verlag.
- Stilman, M., & Kuffner, Jr., J. J. (2007, April). Manipulation planning among movable obstacles. In *To appear in proc. of the IEEE intl. conf. on robotics and automation (ICRA)*. Rome.
- Sung, R., Corney, J., & Clark, D. (2002). Octree based assembly sequence generation. *J. of Computing and Information Science in Engineering*.
- Tanner, H. G., & Kyriakopoulos, K. J. (2000). Nonholonomic motion planning for mobile manipulators. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)* (pp. 1233–1238). San Francisco, CA.
- Teller, S. (1998). Automated urban model acquisition: Project rationale and status. In *Proc. of the image understanding workshop* (pp. 455–462).
- Thoroughman, K. A., & Shadmehr, R. (2000, Oct). Learning of action through adaptive combination of motor primitives. *Nature*, 407, 742–747.
- Thrun, S., Hähnel, D., Ferguson, D., Montemerlo, M., Triebel, R., Burgard, W., et al. (2003). A system for volumetric robotic mapping of abandoned mines. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)*.
- Turk, G., & Levoy, M. (1994, July). Zippered polygon meshes from range images. *Computer Graphics (Proc. of ACM SIGGRAPH)*, 311–318.
- van den Berg, J. P., & Overmars, M. H. (2005, Oct). Roadmap-based motion planning in dynamic environments. *IEEE Trans. on Robotics*, 21(5).
- Wang, C.-C., & Thorpe, C. (2002, May). Simultaneous localization and mapping with detection and tracking of moving objects. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)*. Washington, D.C.
- Yamamoto, Y., Pirjanian, P., Munich, M., DiBernardo, E., Goncalves, L., Ostrowski, J., et al. (2005). Optical sensing for robot perception and localization. In *Proc. of IEEE workshop on advanced robotics and its social impacts* (pp. 14–17).
- Yamauchi, B. (1997, July). A frontier-based approach for autonomous exploration. In *Proc. of the 1997 IEEE intl. symp. on computational intelligence in robotics and automation* (pp. 146–151). Monterey, CA.
- Yang, X. S., & LaValle, S. (2000, April). A framework for planning feedback motion strategies based on a random neighborhood graph. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)* (Vol. 1, pp. 544–549). San Francisco, CA.

Zucker, M., Kuffner, Jr., J. J., & Branicky, M. (2007, April). Multipartite RRTs for rapid replanning in dynamic environments. In *Proc. of the IEEE intl. conf. on robotics and automation (ICRA)*. Rome, Italy.