Spatial Reasoning and Robotics

M.Soutchanski
10 King's College Road SF3303
Dep. of Computer Science,
Univ. Of Toronto,
M5S 1A4 Canada,
email: mes@cs.toronto.edu

May 16, 1996

A paper for the depth oral examination
Friday May 31, 1996
1:30 PM, LP266

This document submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

Contents

1 Introduction **1**

2 Exact information and accurate control **3**
2.1 The complexity of path planning 5
2.2 Complete algorithms for path planning 6
2.3 Incomplete and approximate algorithms for path planning 7
2.4 Applications 8

3 Motion planning under uncertainty **9**

4 Qualitative (topological) spatial information **13**
4.1 Metrically-based topological information 13
4.2 Purely qualitative reasoning about space and time 16

5 Absence of prior information **20**
5.1 Map construction 20
5.2 Skill acquisition 21

6 Future Work **22**

7 Glossary **23**

1 Introduction

This review is by no means a complete account of the literature related to the representation of knowledge for reasoning about actions in space and motion planning.

First, the representation of knowledge about space is the very large area of research, and one can easily count hundreds of papers published in proceedings of several different international conferences. In particular, knowledge about space can be represented for purposes of qualitative reasoning (how spatial arrangements change in time), for geographic systems, for qualitative navigation and other applications. All these applications ultimately require the design of convenient data structures for representing spatial information (accurate or approximate) and managing large spatial databases. This area is partially covered in several recent reviews and books [Davis, 1990, Topaloglou, 1991, Faltings & Struss, 1992, Frank, Campari, & Formentini, 1992, Frank & Campari, 1993, Hernández, 1994, Herring & Egenhofer, 1995].

Second, reasoning about actions is a different area that deals with logical formalizations of action effects and abstract planning of actions in incompletely known environments. Actions can be understood as a high-level description of the ongoing behavior of a robot, but only in a few papers actions have been explicitly interpreted as performed by a robot in a real space. Traditionally, this area of research is more concerned with formal properties of action theories, the planning of activity in an abstract setting and solutions of the frame and ramification problems for deterministic, nondeterministic and compound actions (procedures).

Third, motion planning is a mature and mathematically advanced area of research with a basis in computational geometry, topology, the theory of complexity and differential equations. The motion planning problem can be roughly defined as follows: how can a robot in 3D or 2D space cluttered with obstacles decide (under certain information conditions) what motions to perform in order to achieve goal arrangements of physical objects specified on a high-level language. In many solutions proposed so far the task level specifications [Lozano-Pérez, 1982] are left outside and attention is paid mainly to computing the safe path for a robot (manipulator) amidst obstacles, given a (partially) known geometrical description of a workspace. As in the two other areas mentioned before, motion planning literature is very extensive and includes hundreds of publications. In particular, several research monographs and textbooks discuss different aspects of this problem, among them [Schwartz & Yap, 1987, Schwartz, Sharir, & Hopcroft, 1987, Canny, 1988, Donald, 1989b, Latombe, 1991, Fujimura, 1991, Burger & Bhanu, 1992, Lozano-Pérez *et al.*, 1992, Kramer, 1992].

Several researchers have acknowledged the necessity of a mutually fruitful collaboration among aforementioned areas [Lozano-Pérez, 1982, Hayes, 1985b, Donald, 1989b, Davis, 1990, Forbus, Nielsen, & Faltings, 1991, Cui, Cohn, & Randell, 1993, Kaufman, 1991, Latombe, 1991, Homem de Mello & Lee, 1991, Lozano-Pérez *et al.*, 1992, Brafman, Latombe, & Shoham, 1993, Shanahan, 1995]. The subordination of a motion planning problem to a task planning problem specified in a logical language will allow a planner to decide what geometrical information is necessary and will guide the search of a geometrically and kinematically feasible motion amidst obstacles. However, integrating task planning and motion planning is a very ambitious enterprise and few results have been achieved in this domain. Knowledge-based motion planning is a still emerging area of research with the basis in three vast areas outlined above. Accordingly, my main selection criterion is whether a publication contributed towards achieving the ultimate goal: how to create autonomous robots (equipped with noisy actuators and sensors) capable accept high-level descriptions of tasks and execute them, possibly performing motions in the real world [Nilsson, 1969, Lesperance *et al.*, 1994]. My review has been written with this task in mind. To make my task feasible I disregard those papers which consider only reasoning about abstract actions.

I divide my review into sections in accordance with what information is supposed to be available and how moving bodies are rendered. Thus, I put in one group all publications whose authors consider similar information requirements and levels of abstraction for moving bodies. Whenever possible, I will indicate what publications take into account task level specifications. The main classification schema is the following.

- Accurate numerical (geometric, kinematic and/or dynamic) information is given beforehand and robot actuators are perfect (errors in control are negligibly small).

- There are errors in the initial geometric models, robot control and sensors, but these errors are contained within bounded regions. In other words, the initial and goal position of the robot are not known exactly, and the actual locations of obstacles can be different from those in the robot's model, but all these errors are bounded.

- Only qualitative or topological information is available.

- There is no prior information about the workspace where the robot has to work. It relies on its sensors (which are not perfect) at execution time to obtain the geometric information needed to accomplish its tasks.

In the sequel some of these cases are subdivided into classes which correspond to more refined assumptions about motion. *Italics* is used to emphasize notions defined in the section "Glossary".

2 Exact information and accurate control

> "An AI task planner typically represents a task using logic-oriented constructs (e.g., predicate calculus language) which contain little information about the actual geometry of the objects and the spatial relations among them. Therefore, a plan generated by the task planner may turn out to be infeasible when developed at the motion level."
> [Latombe, 1991]

Because the assumption that exact geometric and dynamic information is available beforehand facilitates theoretical analysis, a majority of results in motion planning have been obtained under this assumption. The following additional assumptions simplify analysis of motion significantly.[1]

Let both moving object $\mathcal{A}$ and obstacles $\mathcal{B}_1, \ldots, \mathcal{B}_k$ occupy some regions in $\mathbf{R}^n$ (n=2,3) and be such that they are

1. solid; i.e., no geometrical point can be inside more than one object;
2. rigid; i.e., the position of every point in the object is characterized by a few parameters;
3. only robot $\mathcal{A}$ is able to move and it can be considered as a uniform object which does not have external links connected by joints;
4. the obstacles $\mathcal{B}_j$ are not movable; i.e., they are stationary and the robot cannot displace them (for example, by grasping or pushing) to other locations;
5. there are no time constraints on how quickly the robot must achieve a goal, and there are no kinematic or dynamic constraints which limit the motions of $\mathcal{A}$ (i.e., $\mathcal{A}$ is a *free-flying object*).

All mentioned assumptions can be relaxed and there are papers which discuss more elaborate motion planning problems. The basic `Findpath` problem [Lozano-Pérez, 1983] is: how to move $\mathcal{A}$ from an initial position and orientation to a goal position and orientation without causing collisions with the $\mathcal{B}_j$. Any *path* τ that specifies a continuous sequence of positions and orientations of $\mathcal{A}$ avoiding contact with $\mathcal{B}_j$'s and leading from the initial to the final position and orientation can be considered as a solution of this problem. Figure (1)

[1] Ideally, they should be stated formally, because reasoning about motion would be more flexible, if it were clear on what assumptions it depends logically.

illustrates the definition of the basic `Findpath` problem for convex polygons. If both $\mathcal{A}$ and $\mathcal{B}_j$ are two-dimensional, then $\mathcal{A}$ has two translational and one rotational degree of freedom. If both a moving body and obstacles occupy regions of $\mathbf{R}^3$, then $\mathcal{A}$ has 3 translational and 3 rotational degrees of freedom. The notion of *configuration* is introduced to unify the treatment of degrees of freedom: the configuration of $\mathcal{A}$ is a set of independent parameters that characterize the position of every point in $\mathcal{A}$.

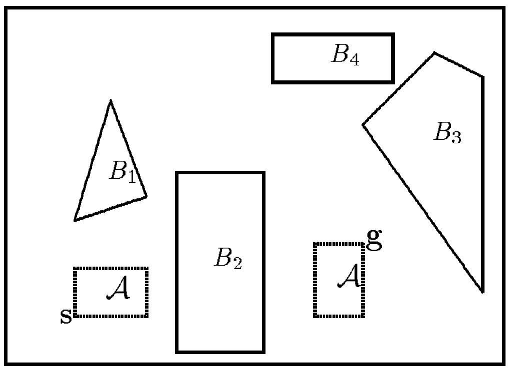

Figure 1: The `Findpath` problem is to find a path for A from s to g that avoids collisions with the B_j .

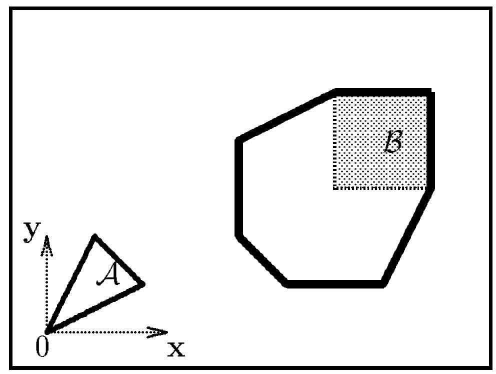

Figure 2: If the point 0 is on or inside the bold face border after a translation of $\mathcal{A}$, then there is an intersection between $\mathcal{A}$ and the square $\mathcal{B}$

The major step in solution of both basic `Findpath` problem and its straightforward extensions is to reduce the task of planning the motion of a dimensioned object to the task of planning the motion of a point. This reduction is attained by introducing the notion of *configuration space* [Udupa, 1977, Lozano-Pérez, 1983, Latombe, 1991, Lozano-Pérez *et al.*, 1992]. The main idea behind this reduction is equivalence between `Findpath` problem and planning of a free motion for a point representing $\mathcal{A}$ in a *configuration space* where obstacles $\mathcal{B}_j$ are expanded by the shape of $\mathcal{A}$. The point which represents $\mathcal{A}$ is the origin 0 of a moving Cartesian frame embedded in $\mathcal{A}$. Figure (2) depicts a *configuration space obstacle* (a territory inside the bold face border). It is a geometric object that represents all the positions of the object $\mathcal{A}$ (assuming $\mathcal{A}$ can translate, but cannot rotate) with respect to a fixed Cartesian frame that cause collisions with the square $\mathcal{B}$. If $\mathcal{A}$ were able to rotate, then the shape of the corresponding configuration space obstacle would be different for each particular orientation of $\mathcal{A}$ with respect to a fixed Cartesian frame. However, if $\mathcal{A}$ is a circle and the origin of a moving Cartesian frame is in its center point, then the circle's orientation is irrelevant and the configuration space remains two-dimensional. The configuration space obstacles for the circle are objects $\mathcal{B}_j$ grown by the circle radius.

Thus, in general case for a body in $\mathbf{R}^2$, the *configuration space* is three-dimensional; if both $\mathcal{A}$ and obstacles are polygons then (after an appropriate non-redundant parameterization) configuration space obstacles are *semi-algebraic sets*. For a body in $\mathbf{R}^3$ the *configuration space* has six dimensions. [Latombe, 1991] describes the differential and topological structure as well as provides the extensive overview of algebraic and geometric properties of these spaces when all objects (including moving body $\mathcal{A}$) are represented by *semi-algebraic sets* (in particular, he analyzes the simplest cases when $\mathcal{A}$ and obstacles are polygons in $\mathbf{R}^2$ or *polyhedral regions* in $\mathbf{R}^3$).

The algorithmic advantage of the notion of configuration space is that the intersection of a point with a set of objects is easier to deal with than the intersection of objects themselves. For this reason, the idea of a *configuration space* is central for motion planning literature. The methodological advantage of this formulation is that it allows us to distinguish those proposals which consider spatial properties and their dynamics on purely qualitative level[2] from those proposals which suppose that a kind of metric (possibly uncertain) information is available[3].

To the best of my knowledge, nothing has been done in the case when the moving body or obstacles do not satisfy requirements (1), (2) or both (see above). Relaxing the assumption (3) gives rise to the straightforward generalization of the basic problem: the case when a robot has several joints or the case when several robots must move independently. In all these cases, the moving system has many degrees of freedom; the appropriately generalized `Findpath` problem is sometimes referred to as the "piano movers' problem", or as the "generalized movers' problem" [Schwartz, Sharir, & Hopcroft, 1987, Schwartz & Yap, 1987, Reif, 1987, Reif & Sharir, 1994]. Another important extension (when the assumption 4 is relaxed) arises when obstacles cannot move themselves, but their motions are under robot's control [Erdmann & Lozano-Pérez, 1987, Hopcroft & Wilfong, 1986]. The presence of movable objects is a serious challenge for a motion planner and demands more sophistication. Instead of reporting a failure in the case when there exists no feasible path to a goal in some given arrangement of the workspace, the planner may be able to create a new path by moving some obstacles to other locations. Hence, a motion plan may appear as a sequence of alternating transit and transfer motions [Wilfong, 1988, Latombe, 1991, Koga *et al.*, 1994]. Finally, without the assumption (5), when we do consider kinematic or dynamic constraints, the job of a motion planner becomes even more complicated [Canny & Reif, 1987, Canny, 1988, Donald *et al.*, 1993, Reif & Sharir, 1994]. Many papers discuss motion planning in the presence of moving obstacles for a robot with constraints on its velocity and acceleration [Canny, Rege, & Reif, 1991, Fujimura, 1991, Donald & Xavier, 1990, Roos & Noltemeier, 1991].

2.1 The complexity of path planning

In this section I review results about (worst case) lower bounds on complexity of the generalized `Findpath` problem. Analysis of motion planning problems is important because it provides complexity bounds for the computational geometry problems in the case when we have exact information about shapes of obstacles. Thus, it reveals difficulty of the problem if we need to compute fine-tuned motions. Moreover, it shows what parameters are crucial and suggest how real robotics problems should be simplified or reformulated to make solutions feasible.

[2] Hence, such proposals have nothing comparable to *configuration space*.

[3] it can be abstracted on a qualitative level, if desired

First, I give the definition of the generalized mover's problem (GMP) which is, roughly, a `Findpath` problem when a robot has several joints (links, bodies). Let a robot in $\mathbf{R}^3$ be represented as a *rational polyhedron* P; such a polyhedron P can be encoded as a finite binary string. Assume also that: (1) P consists of a finite set of (rational) polyhedra $\{p_i\}$, which are freely linked at distinguished linkage vertices $\{v_k\}$; (2) all obstacles are static *rational polyhedra*; (3) given initial and final rational positions of P are *legal*. We may encode this description as a binary string; the size of GMP input is the length n of this encoding. Note that n depends on the number of polyhedra composing P (in other words, on the number of degrees of freedom of P). The GMP is to determine the existence of a legal movement between initial and final positions. The set of strings corresponding to GMPs with a positive answer constitutes a language. Let M be a Turing machine that has a single read/write tape with $s(n)+2$ tape cells, where $s(n) \geq n$ is a space bound which is a function of the input length n; $DSPACE(s)$ denotes the set of languages accepted by a deterministic Turing machine M with the space bound $s = s(n)$. By using the degrees of freedom of P to encode the configuration of a polynomial space-bounded Turing machine[4] M, and by designing a set of obstacles which force the movement of robot P to simulate the computation of M, [Reif, 1987] proved that

Theorem 1 *(Reif, 1979) The generalized mover's problem is $PSPACE$-hard.*

Several subsequent publications reported $PSPACE$-hardness for a variety of path planning problems with simpler or more specific robotic systems [Schwartz, Sharir, & Hopcroft, 1987]. In particular, [Hopcroft, Schwartz, & Sharir, 1984, Hopcroft & Wilfong, 1986] consider the 'Warehouseman's Problem', formulated for several independent rectangular robots translating in an empty rectangular workspace with sides parallel to the sides of the workspace boundary. The problem is to plan the coordinated motion of the rectangles between two given configurations, so that they do not intersect. This problem is $PSPACE$-complete. Thus, lower bounds on the computational complexity of GMP and similar problems show (it is reasonable to believe) that single-exponential time may be required to solve these problems in the worst case. Furthermore, there is an algorithm [Canny, Grigor'ev, & Vorobjov, 1992, Canny, 1993] that needs in the worst case single-exponential time (in the number of degrees of freedom); thus, it establishes a "close" upper bound on the problem complexity. The complexity of this algorithm is discussed below.

2.2 Complete algorithms for path planning

As far as complete[5] algorithms for generalized `Findpath` problem are concerned, several algoritms with different complexity were proposed [Schwartz & Sharir, 1983, Schwartz, Sharir, & Hopcroft, 1987, Canny, 1988]. To plan an obstacle-free motion in the working space, we need to find a path through the configuration space that lies entirely within the *free space*. In many cases, the geometric constraints arising from obstacles have an algebraic description. Consequently, the free space can be defined as a *semi-algebraic* set, which is a subset of some

[4] Thus, for such machines $s(n)$ is a polynomial function of n.

[5] In the sense that they find a legal path between two specified configurations if it exists.

real space $\mathbf{R}^f$ formed by intersections and unions of algebraic surfaces.[6] The dimension f of this space depends only on the number of degrees of freedom. Thus, obstacles can be represented by sentences with no more than f variables in the *Tarski theory of reals*.

Several good sequential and parallel algorithms are available for deciding the theory of reals and for elimination of quantifiers over real variables [Collins, 1975, Grigorev, 1988, Grigorev & Vorobjov, 1992, Heintz, Roy, & Solernó, 1989, Renegar, 1992]. One of them, based on cellular algebraic decomposition [Collins, 1975, Ben-Or, Kozen, & Reif, 1986], is used to decompose the free space, and the existence of a collision-free path is demonstrated by computing cell adjacency [Schwartz & Sharir, 1983]. However, this algorithm runs in twice-exponential time in f, the dimension of the configuration space (because algebraic cell decomposition requires twice-exponential time). If dimension of a configuration space fixed, and the free space is a set defined by n polynomial constraints of maximal degree d, a collision-free path is computed in time polynomial in both n ('geometric complexity') and d ('algebraic complexity').

A different approach is taken in [Canny, 1988], based on the concept of a road-map, which is a one-dimensional subset of the robot's configuration space. Let $S \subset \mathbf{R}^f$ denote a subset of the free space. J.Canny defines a road-map which has the property that any connected component of S contains a single connected component of the road-map. Starting from an arbitrary point p that belongs to S, it is possible to construct a path from p to a point on the road-map. Thus, given any two points in S, we can determine whether they lie in the same connected component of S, and, if they do, a candidate path can be found between them. In addition to proving a key property of connectivity in motion planning, the author also proves that the road-map is a semi-algebraic set and derives a representation for the road-map from S. [Canny, 1988] proposes an algorithm (under some assumptions) for finding paths that has complexity $(n^f \log n) \cdot d^{O(f^2)} w^2$, where n is the number of polynomials, d is their maximum degree, w is the maximum length of their coefficients in bits, and f is the number of variables; thus his algorithm needs single-exponential time in f. Because the lower bound on the number of components is $\Omega((nd)^f)$, his algorithm is nearly optimal in terms of n. [Canny, Grigor'ev, & Vorobjov, 1992, Canny, 1993] later improved the original road-map algorithm (making it applicable in the general case).

Many other (complete) algorithms have been proposed for the solution of the basic `Findpath` problem and its variations in other settings [Brooks, 1983, Schwartz & Yap, 1987, Donald, 1987, Latombe, 1991, Reif & Storer, 1994]. Some address simplified, but still realistic problems, and solve them with a reasonable efficiency.

2.3 Incomplete and approximate algorithms for path planning

There are also several approximate methods [Lozano-Pérez, 1981, Lozano-Pérez, 1982, Brooks & Lozano-Pérez, 1985, Kambhampati & Davis, 1986]. They are not guaranteed to find a free path if one exists, but because the precision of the approximation can be tuned, these methods are said to be "resolution-complete". For example, algorithms known as approximate cell decomposition methods divide the working space into cells of a simple shape (e.g., rectangloids). Interiors of some cells may lie inside the free space, while other cells may occupy

[6] An algebraic surface is the set of points in $\mathbf{R}^f$ that satisfy a polynomial inequality.

part of the *free space* and part of a configuration obstacle. These border-line cells can be recursively decomposed until a resolution limit is achieved or there are no border-line cells. At a certain level of resolution, the connectivity graph can be easily constructed from information about adjacency only between cells inside the free space. Thus, algorithms of this type employ a special data structures for the representation of spatial properties known as 'quadtrees'[7] and 'octrees'[8] [Samet, 1990b, Samet, 1990a]. The navigation program of SRI's robot Shakey was one of the first attempts to use quadtrees for motion planning [Nilsson, 1969] (the lowest resolution available for Shakey was 12 inches). Approximate methods are also used for real-time collision avoidance [Shaffer, 1992] and for kinematic reasoning [Olivier, Ormsby, & Keiichi, 1995]. E.Davis considers a different approach to the issues of shape approximation in kinematic mechanical systems and configuration space approximation [Davis, 1995b, Davis, 1995a].

As far as incomplete[9] algorithms for the generalized `Findpath` problem are concerned, the most popular method is based on a collision-avoiding attractive *potential field* over the configuration space [Khatib, 1986, Latombe, 1991]. The idea is to apply an artificial force to the mobile robot so that it is repelled by the obstacles while being attracted to the destination point. In comparison to other methods, potential field methods can be very efficient because they are essentially fastest descent optimization methods. The three major problems of this approach are that the robot may get trapped in local minima, that the robot's path may oscillate in certain situations and that some common robotic devices (such as car-like vehicles) cannot move easily in the way that the potential field is directing them. There are several proposals on how to improve the *potential field* method: (1) design better potential functions [Rimon & Koditschek, 1992]; (2) mingle the basic *potential field* method with a global motion planning method or another mechanism to escape from local minima [Canny & Lin, 1990, Barraquand & Latombe, 1991, Qin, Cameron, & McLean, 1995].

2.4 Applications

> "Surprisingly, the robot program to assemble a typewriter is not likely to have any subroutines that would be useful in assembling a bicycle or even a different model of typewriter. (...) If robots are to fulfill their scientific and economic potential, this programming bottleneck must be eliminated."
> [Lozano-Pérez *et al.*, 1992]

Because precise geometric information can be expected either in computer graphics (automatic animation) or in computer-aided design and computer-aided manufacturing systems (assembly planning), algorithms proposed for solution of the generalized `Findpath` problem may have significant impact [Kusiak, 1990, Popplestone, Liu, & Weiss, 1990, Homem de Mello & Lee, 1991, Lozano-Pérez *et al.*, 1992, Gottschlich, Ramos, & Lyons, 1994, Wilson & Latombe, 1994]. Meanwhile, the importance of exact geometric information should not be overestimated. Imagine that an animation generator [Koga *et al.*, 1994, Bergacker, Williams,

[7] A tree such that each vertex represents a rectangle in $\mathbf{R}^2$ decomposed into 4 identical rectangles.

[8] A tree such that each vertex represents a cube in $\mathbf{R}^3$ decomposed into 8 identical rectangles.

[9] In the sense that they can miss a legal path even if one really exists.

& Kalita, 1994] has to plan the motion of an agent in a digital garden containing simulated bushes, trees and flowers from one point on a lawn to another. It seems that the easiest way to treat this scene would be to consider grass in this digital garden as a smooth surface and localize the `Findpath` problem to the smallest region that includes departure and destination points. Other task level properties relevant to the scene may determine how to approximate shapes of objects and guide decisions whether collision with an object may or may not occur.

The authors of [Lozano-Pérez *et al.*, 1992] describe the system (called `Handey` after 'hand+eye') developed in MIT. The research was centered around interrelationships between task-level specifications and implemention of commands on the motion level. The book reports the results of several years efforts towards developing a general-purpose model-based industrial robot capable of picking up user-specified objects and placing them at user-specified positions (these manipulation problems are called pick-and-place problems). `Handey` is a robot programming system with a substantial level of pick-and-place competence. It means that `Handey` needs a description of the task rather than a specification of the robot motions required to carry out the task. Given a goal to achieve and a polyhedral description of the environment and of the robot, `Handey` computes a desired sequence of commands (with numerical parameters) to carry out the task. The key property of a task-level robotic system is that it works reliably in different environments. There are many examples in the book which show that motions planned initially on purely local considerations may happen to be infeasible during the execution. For example, obstacles in a position where an object has to be grasped constrain possible grasping points, but when an arm with the object approach a destination, it may happen that due to obstacles there, arm cannot place the object without a collision. In those cases when the task has a solution, robotic system should be able to find a temporal position where the object has be regrasped in such a way that it can be placed in the destination. Thus, regrasp planning becomes an important issue related to recovery from a failed motion sequence. Because there are infinitely many qualitatively different arrangements of obstacles in the destination, the robot's task cannot be solved using a standard library of reactive procedures. Authors of the book conclude that "the key problem in grasp planning is how to choose grasps that enable successful completion of the whole task, not just the initial grasp." They also points out that task-level programming in workspace with substantial uncertainty still requires fundamental research.

3 Motion planning under uncertainty

> "Two difficulties became apparent with Shakey. The first was that uncertainty mattered. Someone had to implement low-level routines to deal with uncertainty. The second difficulty was that the symbolic states used to represent the robot's configuration were too simple. Once one proceeds to more complicated tasks, such as the assembly of intricately shaped tasks, these symbolic states are insufficient. Instead, it seems to be necessary to consider geometry in detail".
> [Erdmann, 1995]

Geometry was discussed in the previous section. The second thread of work (that deals with uncertainty) is discussed in this section.

Robot motion planning with uncertainty is a sufficiently explored area [Lozano-Pérez, Mason, & Taylor, 1984, Canny & Reif, 1987, Erdmann, 1986, Donald, 1989a, Latombe, Lazanas, & Shekhar, 1991]. The uncertainty can arise from several sources, including (partial) information about the location of some obstacles, sloppy control, and imperfect sensors. Sensing and control uncertainty can be modelled either in probabilistic terms or in terms of "adversarial" influences. The latter distort sensing and control in a systematic but unpredictable way. For example, it is assumed that the actual velocity of a robot is a vector in a cone centered around a commanded velocity (and the angle of this control uncertainty cone is acute) such that orientation and the modulus of the actual velocity can vary arbitrary (but continuously) within these finite uncertainty bounds. In regard to the sensing error, it is further assumed that the difference between an actual (unknown) configuration q_a and the sensed configuration q varies from one measurement to another. The only assumption about uncertainty made by the "adversarial" model is that variations are bounded: $\| q_a - q \| \leq \epsilon$. Consequently, it is **not** assumed that the average value $\frac{\Sigma_{i=1}^{n} q_i}{n}$ of repetitive measurements converges (in a probabilistic sense) to the actual value q_a if n increases ("adversarial" influences can be considered as a multiplicative noise). In contrast, modelling uncertainty in probabilistic terms would assume that identification of the actual configuration q_a is possible with some statistical degree of confidence. Modeling control uncertainty in probabilistic terms can give rise to stochastic versions of the *potential field* method where motion follows the stochastic quasi-gradient of a potential over *configuration space* and new directions of motion are decided upon measuring values of a potential. Under certain assumptions,[10] stochastic quasi-gradient algorithms are able to find the minimum value of a function (possibly, measured with noise). Hence, methods developed in stochastic optimisation theory could be applied to successful tracking of a goal in this setting. However, the applicability of this approach seems questionable because stochastic optimisation algorithms usually have relatively low speed of convergence. In the section 5 I briefly review some other approaches based on different ideas from the theory of learning.

Almost all publications consider the "adversarial" model of uncertainty. The archetypical problem considered in these publications is the 'peg-into-hole' problem. More realistic assembly problems (3D objects must be mated) can be viewed as a point-into-volume problem in some appropriate configuration space [Lozano-Pérez, 1981, Lozano-Pérez, 1983]. The generalized `Findpath` problem with uncertainty is to produce a plan which is guaranteed to succeed even if the robot cannot perfectly execute it due to control error. Early papers consider the enhancement of planning and execution systems by special feedback loops and force control strategies that combine motion and sensing commands in a manner that reduces uncertainty. One of the strategies, compliant motion, was developed for robots equipped both with a position and with force sensors. Compliant motion [Mason, 1981] occurs when a robot is commanded to move into an obstacle, but instead of pursuing a command, it complies with the geometry of the obstacle. In contrast to the ordinary `Findpath` problem, which can be addressed without regard to the robot dynamics, compliant motion was studied using certain dynamic models that describe the behavior of robots in contact with obstacles: 'generalized spring' and 'generalized damper' models [Mason, 1981, Lozano-Pérez, Mason, & Taylor, 1984,

[10] Typical assumptions are the following: the expected value of noise is 0; its variance is bounded by a known amount and the potential function is differentiable on configuration space.

Khatib, Craig, & Lozano-Pérez, 1992, Latombe, 1991]. They are taken into account by motion planners because compliant motion strategies employ obstacles as a kind of landmarks to compensate for uncertainties in control and lead robots to the goal reliably.

Early work on planning in the presence of uncertainty is based on the idea of skeleton strategies [Brooks, 1982]. Skeleton strategies are programs converted from the task specification by a global motion planner assuming ideal position information and control accuracy. The strategies obtained after the first phase contain feedback parameters because they should be tailored to fit the particular geometric environment. During the second phase, skeleton strategies are iteratively modified using sensor-based readings and sensor-guided motions. [Lozano-Pérez & Winston, 1977] describe a task-level planner called LAMA which uses geometric simulation to predict the outcomes of plans. These outcomes are used to constrain variable values in skeleton plans. Thus, the skeleton refining approach deals with imperfections of the data and control errors only in a second phase of planning. Consequently, this approach can be successful only if uncertainty does not affect the overall structure of the skeleton plan generated initially.

[Lozano-Pérez, Mason, & Taylor, 1984] propose formal method of planning with uncertainty based on *preimages*. Their method directly incorporates the effect of uncertainty within the planning process. The preimage of a goal is the set of positions in a *free space* from which particular commanded motions are guaranteed recognizably to attain their goal, despite errors in control and uncertainty in sensing. Given bounds on uncertainties, the maximal *preimage* is the biggest area from which the robot can start. Typical motion commands include "move in direction θ with speed v". The method (called preimage backchaining) consists of iteratively computing preimages of the goal, preimages of computed preimages being taken as intermediate goals, for various possible motion commands. This recursive computation succeeds eventually if some preimage contains the initial subset of configurations in which the robot is known to be before motion starts. Thus, a chain of preimages determined by a planner specifies a sequence of motion commands guaranteed to reach the goal. The sequence of commands is the desired motion strategy. The preimage backchaining approach is applicable even when there is no uncertainty and with different dynamic models of contact. In fact, it corresponds to a very general planning method known as "goal regression" [Waldinger, 1977, Nilsson, 1980]. The geometric specificity manifests itself in representation of the search space and in the way how particular regressions (preimages) are computed.

The preimage approach provides a formal framework for approaching the problem of planning with uncertainty, but raises difficult computational issues. Several subsequent papers propose practical methods for computing preimages in low dimensional *configuration spaces*. [Erdmann, 1986] notes that it is worthwhile to separate the issue of goal reachability and the issue of recognizing that goal is attained. He suggests replacement of a goal region by one or more subsets such that their attainment is recognizable. After that, the problem reduces to planning of motions that reach one of these sub-regions. The paper provides a method for constructing regions called backprojections, from which particular motions are guaranteed to reach a goal. A backprojection of a goal differs from a preimage by the fact that the condition of recognizability is ommitted. Hence, the backprojection of G for any given motion command is an upper bound on the maximal preimage (if any) of G for the same command, over all possible termination conditions. [Latombe, Lazanas, & Shekhar, 1991, Lazanas, 1994] develop other approaches for computing preimages, proceeding from Erdmann's idea of backprojection. They also describe an algorithm for the generation of conditional motion strategies.

[Donald, 1988, Donald, 1989b] introduces the notion of model error, which is initial uncertainty in the geometric models of the environment and of the robot. In particular, he shows that model error can be dealt with by adding a dimension to configuration space for every parameter in the workspace whose value is not known accurately. For example, if the initial x coordinate is supposed to be somewhere in (l_1, l_2), then the 'generalized configuration space' will have two dimensions instead of one: x and δx , such that values of δx vary in (l_1, l_2). Donald also moves away from the requirement that a strategy, to be considered legitimate, must actually be guaranteed to solve a task in a finite prescribed number of steps. He introduces the notion of error detection and recovery (EDR) strategy: a motion plan which achieves the goal when it is recognizably reachable and signals failure when it is not. Furthermore, EDR strategies should not be prematurely terminated as failures without strong evidence. EDR strategies are more general than guaranteed strategies, in that they allow plans to fail.

A system for navigation in a real-world environment with uncertainty is described in [Cimatti *et al.*, 1992, Traverso *et al.*, 1994]. Its main feature is the ability to integrate goal-directed planning with information acquisition from the external world, and with reaction to failure. [Giunchiglia, Spalazzi, & Traverso, 1994] develop a language based on dynamic logic suitable for reasoning about failures of complex procedures and define a declarative formal semantics for this language.

[Erdmann, 1993] develops a randomization approach to overcome partially the problem posed by Donald: how to reach a goal if success cannot be guaranteed. Randomization accomplishes a task blurring the significance of unmodeled or uncertain parameters probabilistically. For example, using only a sensor that signals goal attainment, the 'peg-into-hole' problem can be solved by random shaking. [Erdmann, 1995] addresses the problem of designing sensors specially suited to a task given a task specification, the robot's action and the uncertainty in control. The key idea is to generate a motion plan[11] by using a backchaining planner that assumes 'perfect sensing'. By 'perfect sensing' he refers to the ability of the sensor to provide information required for the plan to function correctly. For example, a sensor may be able to report for each action whether it makes progress towards the goal. Erdmann argues that "Sensors should not recognize states; sensors should recognize actions", because the robot needs to know only that a certain action will make progress toward accomplishing the task. The robot does not need to know precisely the state of the environment. A designer should try to implement the specification of an abstract sensor by building physical sensors that provide the information described by the progress cones. If no such implementation exists, the designer should consider information-gathering actions to establish this information.

[Brafman, Latombe, & Shoham, 1993] considers one particular problem: given a motion planning instance and a motion direction D, does there exist a sound and complete termination condition? Roughly speaking, a termination condition (i.e., the boolean function on

[11]The plan has to be conditional on sensing because actions may have errors.

the history of possible readings) is sound if at the very first moment when the termination condition evaluates to true, the robot's trajectory is inside the goal region. The termination condition is complete if for every consistent trajectory (trajectory inside error bounds) there is a moment of time such that termination condition at this moment evaluates to true. For a specific class of simple motion plans in $\mathbf{R}^n$, the paper provides necessary and sufficient conditions for the existence of sound and complete termination conditions. The authors define a multi-modal language to formulate general conditions for the existence of sound and complete termination conditions for a broader class of motion plans.

4 Qualitative (topological) spatial information

The representation of spatial information on a qualitative level is the vast area of research with applications in spatial data bases and geographic information systems. There are several reviews [Egenhofer & J.R.Herring, 1990, Topaloglou, 1991, Hernández, 1994, Hernández & Mukerjee, 1995] covering this area and many conference proceedings collect relevant papers [Anger, Guesgen, & van Benthem, 1993, Anger & Loganantharaj, 1994, Anger, Guesgen, & Ligozat, 1995, Frank & Campari, 1993, Herring & Egenhofer, 1995]. I focus only on publications dealing with motion and/or changes of spatial arrangments in time, as I explained in the introduction.

All relevant proposals can be roughly classified into 2 categories. The first one comprises papers where it is assumed that topological relations are ultimately grounded in metrical information. Authors of these papers mention how qualitative relations are derived or abstracted from numerical spatial data. The second category comprises papers whose authors do not mention explicitly how their qualitative relations (having a spatial 'flavour') can be related to metrical information. The two subsections below review these two categories. My notes are fragmentary because, to the best of my knowledge, there is no generaly recognized paradigm in this area of research.

4.1 Metrically-based topological information

> "We claim that there is no purely qualitative, general-purpose kinematics. Unlike qualitative dynamics, where weak representations of time-varying differential equations suffice for a broad spectrum of inferences, weak spatial representations appear virtually useless."
> [Forbus, Nielsen, & Faltings, 1987]

The most firmly established and comprehensive research direction in this category is qualitative kinematics of mechanisms (this direction is usually considered a part of qualitative physics). As we have seen, *configuration space* of multi-body systems is highly dimensional and, as a consequence, planning routes for the simultaneous motion of several bodies is a computationally intractable problem. For example, the configuration space for only two solid and rigid bodies capable of performing arbitrary motions will have $6 \cdot 6 = 36$ dimensions ! Apparently, it is so complicated that it is not possible to plan motions in this space even ignoring kinematic and dynamic constraints on motion. However, in the early 1980s it was realized that in many real mechanisms sub-parts of an assembly perform only a few types of simple motions. In some mechanisms (e.g, clocks), a sub-part moves either along a predefined axis (hence, this part has 1 translational degree of freedom), or rotates around one fixed axis (having only 1 rotational degree of freedom). Following a tradition in mechanics, a pair of interacting objects is called a *kinematic pair*; the whole mechanism is considered as a *kinematic chain* used for transmission forces and motions. Because parts of kinematic pair have only a few degrees of freedom with respect to each other, the configuration space for them is low-dimensional and its analysis is computationally feasible.

Meanwhile, research following [Forbus, 1983] and other seminal papers (reviewed, for example, in [Davis, 1990]) demonstrate that useful spatial representations consist of two main parts: a metric diagram, which includes quantitative information (and thus provides a substrate for numerical, algebraic or bitmap processing – whatever you like), and a *place vocabulary*, which makes explicit qualitative distinctions in shape and space relevant to the current task. A metric diagram can be constructed from configuration place; it contains both symbolic and quantitative information used as an oracle for simple spatial queries. In contrast, a *place vocabulary* (or *region diagram*) is a purely symbolic description of shape and space, grounded in the metric diagram. Each element of this symbolic diagram denotes some contiguous region of space where some important property remains constant. Thus, *place vocabulary* is a graph in which nodes denote regions of space and edges denote an adjacency relation between the regions.

When the ideas mentioned above were associated with each other [Faltings, 1987, Forbus, Nielsen, & Faltings, 1987, Gelsey, 1989, Joskowicz, 1987], it was understood that a combination of quantitative metric diagram and qualitative place vocabulary may support kinematic reasoning: (1) computing potential connectivity relationships between moving objects; (2) computing tessalation of the space of possible object positions into states where connectivity (and some dynamic properties) remain constant; (3) computing all possible transitions between states. Because any real motion of bodies corresponds to a traverse along edges of the resulting graph, it can be used for the prediction of the long-term behavior of a mechanical system. Following this research direction, classification schemas for all existing mechanisms have been proposed [Gelsey, 1989, Joskowicz & Sacks, 1991a, Joskowicz & Sacks, 1991b, Nielsen, 1988] and complicated mechanisms have been successfully analyzed [Faltings, Baechler, & Primus, 1989, Faltings, 1990, Forbus, Nielsen, & Faltings, 1991, Faltings, 1992].

The major drawback of 'qualitative mechanics' is hidden in the excessive amount of metrical information on which it depends. As [Davis, 1988, Kaufman, 1991] note, if one wants to consider behavior of a mechanism with a slightly different geometry (say, one extra teeth on a wheel), the corresponding configuration space must be computed again and the new place vocabulary can be completely different. It is not obvious how region diagrams can be approximated. Moreover, it is not clear how general knowledge about a domain can guide the construction of a region diagram and provide hints what parts of a diagram can be safely ignored. For purposes of robotics and navigation planning, it seems important to learn: (1) the basic idea that topological information should be related to configuration space; and (2) techniques how reasoning about motion in spatially distributed systems can be implemented.

It seems that ideas from qualitative kinematics will be applicable to navigation planning in relatively uncluttered environments, only if original shapes can be approximated. However, even for simple shapes, qualitative kinematics-based approach is very inefficient if qualitative graph is monstrously large and we should follow extremly long chain of transitions along edges of the graph to predict a result of motion. Examples of this problem are provided in [Davis, 1988, Kaufman, 1991] and are considered in the next sub-section.

Several other papers consider metrically-based qualitative spatial reasoning: [Funt, 1980, Kuipers & Levitt, 1988, Kuipers, 1990, Ananthanarayanan, Goldenberg, & Mylopoulos, 1994, Rajagopalan, 1994a, Rajagopalan, 1994b, Pang & Trudel, 1995, Gagné, Pang, & Trudel, 1995, Faltings, 1995]. For example, [Faltings, 1995] attempts solve the `Findpath` problem using only qualitative information about overlap between regions. The goal is to avoid algebraic representations and the complex computations they involve. The prototype of his program demonstrates the topological reasoning techniques for 2D objects. The input to the program is given in the form of three collections of convex bitmaps (i.e., occupancy grids with fixed resolution). These three collections represent fixed obstacles, the moving body and "bubbles", respectively, where "bubbles" are regions covering free space, but possibly overlapping with obstacles and each other. His algorithm preprocess all bitmaps to determine all simultaneous overlaps between triples of regions. Information about those overlaps is employed for finding a qualitative path between the initial and final placements of the moving body. Unfortunately, the details are left outside of the scopes of his paper. It remains unclear on what principles he selects regions covering free space and how the algorithm backtracks, if it finds that the current resolution is too coarse. He also does not demonstrate whether his bitmaps-based qualitative approach performs better than other approximate methods for the solution of the `Findpath` problem.

A point based spatio-temporal logic is developed in the sequence of papers [Pang & Trudel, 1994, Pang & Trudel, 1995, Gagné, Pang, & Trudel, 1995]. The authors consider their work as an improvement over previous proposals desribed in [Davis, 1990, Kuipers & Levitt, 1988, Kuipers, 1990]. The intended domain is a two-dimensional world of horizontal and vertical lines, traffic lights and robots delivering goods. These papers describe a language, the interpretation and show a few examples how formalizations work. The proposed formalism is a multi-sorted dialect of the first order logic. Fluents and actions are parametrized with one temporal argument. Because the authors assume that the robot travels only along predefined (and perfectly known) corridors (represented by straight lines) without any errors in sensing or control, the optimal path finding problem reduces to search in a graph with arcs labeled by distances. They assume that motion of robots is controlled by traffic lights (to prevent collisions), but in their formalization lights change colors periodically, regardless whether any robot recently passed through a corridor. Unfortunately, the frame and ramification problems are left unrecognized in these papers. As a consequence, there is no way to prove whether particular traffic and scheduling laws will prevent collisions and/or optimize delivery time.

[Bäckström, 1990] represents geometrical knowledge in terms of logical theory for the purposes of modeling mechanical assembly processes. He considers axi-symmetric objects with rectangular projections on the (x, y)-plane and calls this model the $2D+$ world. Objects can have (rectangular) holes inside. Whenever two or three objects are assembled by inserting one of them in a hole inside another, they form a new composite object. Bäckström introduces several function symbols and 'primitive' predicates (assuming for simplicity that all objects have either horizontal or vertical orientation) and defines a set of 'derivable' predicates in terms of 'primitive' predicates. For example, to express overlap between objects o_1 and o_2 he defines:

$$Overlap(o_1, o_2) \leftrightarrow \exists p[Pinside(o_1, p) \wedge Spinside(o_2, p) \vee Spinside(o_1, p) \wedge Pinside(o_2, p)],$$

where $Pinside(o, p)$ is 'derivable' predicate true iff point p is inside object o, and $Spinside(o, p)$ is 'derivable' predicate true iff point p is strictly inside object o. These 'derivable' predicates are defined by means of obvious relations between the x and y coordinates of a point and those of an object. After a thorough analysis of all possible contact cases, he defines what it means for objects o_1 and o_2 to be in contact with each other at a point p. The set of actions include 4 translational movements (left, right, up, down), 2 rotations (90^o clockwise and 90^o anti-clockwise) and actions $attach(b_1, b_2)$ and $detach(b_1, b_2)$ over bodies. The author proposes using dynamic logic for specifying the results of motion and uses a non-monotonic logic to solve the frame and ramification problems. To simplify the theory, he makes several assumptions. For example, in his $2D+$ world, objects do not slide; consequently, if one rectangle lies on top of another and the lower rectangle rotates, than the higher rectangle rotates too (as if we saw them from the upper viewpoint as projections of nearby blocks). Finally, the author gives a sketch of an algorithm for deductive verifying the correctness of an assembly plan using his logical theory. This paper (written in 1988) is interesting because it is one of the few attempts to include geometric constraints in considering the effects of action. [Davis, 1988, Davis, 1994, Shanahan, 1995] provide alternative approaches.

4.2 Purely qualitative reasoning about space and time

> "(...)The basic issue is that our intuitive local space is, indeed, probably not a topological space. It is certainly not three-dimensional Cartesian space, which contains such wildly implausible objects as space-filling curves and the Alexander Horned Sphere. Many mathematical intuitions at the basis of geometry and real analysis (from which topology is abstraction) seem to be at odds with the way we think about everyday space."
> [Hayes, 1985b]

There are several papers which discuss the evolution of qualitative/topological properties in time: [Cui, Cohn, & Randell, 1993, Cui, Cohn, & Randell, 1992, Cohn *et al.*, 1994, Egenhofer, 1989, Egenhofer & Al-Taha, 1992, Asher & Vieu, 1995, Vieu, 1993, Pianesi & Varzi, 1994, Mukerjee, Manish, & Praveen, 1995].

These researchers consider only qualitative, or topological properties of spatial regions. Several papers rely upon Clark's calculus of individuals [Clarke, 1981, Clarke, 1985], or on similar ideas of basic "interval"-type relations between regions. Despite that original intuition can come from different sources (topology, mereology and so on), the resulting formalism share many important features. As an example, I consider one school of research. In a series of papers, [Cui, Cohn, & Randell, 1992, Randell & Cohn, 1992, Cui, Cohn, &

Randell, 1993] develop an axiomatic approach towards the description of spatial properties; in particular, they introduce a set of definitions and consider algorithms for reasoning about changes in spatial arrangments. They start with one symmetric and reflexive binary relation $C(x, y)$ interpreted as "x is connected with y" defined on regions x, y. After that, they define several other relations:

$DC(x, y) \rightleftharpoons_{def} \neg C(x, y)$ (x is disconnected from y);

$P(x, y) \rightleftharpoons_{def} \forall z[C(z, x) \rightarrow C(z, y)]$ (x is a part of y);

$PP(x, y) \rightleftharpoons_{def} P(x, y) \wedge \neg P(y, x)$ (x is a proper part of y);

$x = y \rightleftharpoons_{def} P(x, y) \wedge P(y, x)$ (x is identical with y);

$O(x, y) \rightleftharpoons_{def} \exists z[P(z, x) \wedge P(z, y)]$ (x overlaps y);

$PO(x, y) \rightleftharpoons_{def} O(x, y) \wedge \neg P(x, y) \wedge \neg P(y, x)$ (x partially overlaps y);

$DR(x, y) \rightleftharpoons_{def} \neg O(x, y)$ (x is disjoint from y);

$EC(x, y) \rightleftharpoons_{def} C(x, y) \wedge \neg O(x, y)$ (x is externally connected with y);

$TPP(x, y) \rightleftharpoons_{def} PP(x, y) \wedge \exists z[EC(z, x) \wedge EC(z, y)]$ (x is a tangential proper part of y);

$NTPP(x, y) \rightleftharpoons_{def} PP(x, y) \wedge \neg\exists z[EC(z, x) \wedge EC(z, y)]$ (x is a nontangential proper part of y)

All these basic relations as well as direct transitions from one connectivity state to another are represented pictorially in the figure 3. I call formulas composed from aforementioned predicates RCC–formulas, the set of the definitions above RCC–axioms; respectively, the RCC–theory is the set of formulas derivable from RCC–axioms using standard logical rules of inference. Besides basic relations, the authors introduce also operations on regions and define more complex relations between them.

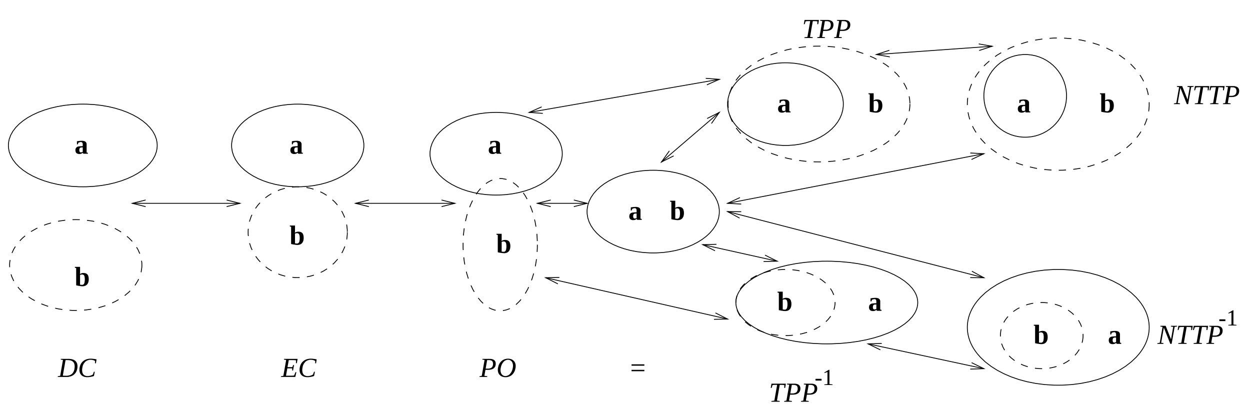

Figure 3: A pictorial representation of basic relations between regions and direct transitions between relations.

Following [Hayes, 1979, Hayes, 1985a] and the 'naive physics programme', [Randell & Cohn, 1992] attempt to axiomatize topological and geometrical issues inherent in modeling of a force pump. In [Cui, Cohn, & Randell, 1992], the authors describe an envisonment-based simulation program. They assume that topological information is extracted from the modelled domain and is expressed in the theory as a set of relation holding between objects.[12] For the temporal part of the theory, the authors consider intervals of time with the standard set of 13 possible relations between them (similar in spirit to relations mentioned above). They represent states as conjunctions of ground atomic formulas corresponding to a particular arrangment of regions. Two different states are characterized by at least one different relation between regions. Given a complete description of the initial state (all relations between regions in the initial period of time) and local envisioning axioms (depicted in figure 3), their program generates an envisionment to predict all possible evolutions of the system.

The drawback of all qualitative formalizations of space based on the primitive concept "connection" is that the resulting theories are too weak to be applicable for robot navigation or for similar domains.

First, these theories are first-order theories (in the logical sense) which may have interpretations without any spatial 'flavour'. For example, it is straightforward to check that if one interprets regions as regular languages over a finite alphabet[13] and stipulate that languages $\mathcal{L}_1$ and $\mathcal{L}_2$ are connected iff their intersection as sets of words is non-empty, then all basic definitions above remain consistent. To get rid of undesirable models, the RCC-theory has to be approached as an interpreted theory in the sense that only certain classes of models are considered legal. However, the example above shows that if one wants to define dimensions, orientations and some other important spatial features in an extension of RCC-theory, then new predicate or function symbols with an intended interpretation should be introduced [Cui, Cohn, & Randell, 1993]. The problem with this approach seems to be in the lack of coherent intuition behind syntactic interrelationships between numerious primitive notions which are introduced. Indeed, rather than to be defined by axioms from an intuitively clear spatial representation, dimensions, orientations and other features will be addressed only on semantical level. It seems that another approach which would be able to encompass the appropriate spatial intuitions explicitly on the syntactic level would be preferable.

Second, [Rogers, 1956] proves that the theory of symmetric, reflexive relations is undecidable (see also [Ershov *et al.*, 1965]). Because, all basic relations are defined in this theory, starting with $C(x, y)$, we find that the RCC–theory is undecidable.

Third, RCC–axioms are epistemologically weak, because from basic relations alone we cannot conclude whether there is a path in a given scene. For example, consider figure 1 and assume that we have a set of RCC–formulas describing the scene on this picture. Obviously, the same set of RCC–formulas will remain true if obstacles B_1-B_4 will expand and change shapes in such a way, that they do not touch each other but all routes between initial and final configurations of $\mathcal{A}$ cease to exist. Thus, the solution of the `Findpath` problem is

[12] It is assumed that at every moment in time, each object occupies a unique region.

[13] To comply with extensions of RCC–theory which admit both non-atomic regions and points, non-point regions should correspond to regular languages defined by regular expression with Kleene *-operator, while points have to be interpreted by languages defined without *-operator.

not invariant to some of those scaling transformations (or shape deformations) over regions which preserve RCC-relations between these regions. Therefore, RCC-theory alone does not provide useful representation of spatial properties for the purposes of navigation planning.

Several other papers discuss logics for representation and reasoning about space, time and motion on a qualitative level: [Davis, 1988, Kaufman, 1991, Egenhofer & Al-Taha, 1992, Galton, 1993, Galton, 1995, Myers & Wilkins, 1995].

[Davis, 1988] proposes a framework for commonsense prediction of solid object behaviour. He considers the following example. Suppose that one releases a small object (die, block or shere) inside orifice of a vertically positioned funnel. The problem is how to prove that the obect will fall through the bottom orifice (assuming that diameter of the funnel is everywhere greater than the size of the object). If one drops a cube inside a funnel and wants to determine which face ends up, then tracing the progress of the cube through all qualitative states is inevitable. Davis points out later: if all we wish to determine is that the cube falls through the funnel, then all these intermediate states do not matter and it should not be necessary to trace them. He writes: "To determine by any kind of simulation that the die falls through the funnel will be almost as difficult as to determine how it lies. This seems implausible." To solve the problem, Davis proposes a first-order axiomatization which embraces geometrical, temporal and physical properties. He argues using the idea of dissipation of energy that an isolated convex object must exit from the bottom of the funnel. Among the major problem encountered during his work, Davis mentiones the representation issue: what can constitute a general geometric description capable of representing arbitrary shapes, position, and physically possible motion.

[Kaufman, 1991] describes an approach to qualitative reasoning grounded in the axiomatization of tolerance spaces (which are weaker than topological spaces); these were first suggested for use in AI by [Hayes, 1985b]. The intuition behind the use of such spaces as models of the physical world is that space has a definite granularity, i.e., it does not make sense to talk about positions at a smaller scale. Because of granularity, all objects occupying regions in a tolerance space can be decomposed into sequences. From this observation, Kaufmann concludes that mathematical induction can be employed in reasoning about the behaviour of physical systems, if a connection between adjacent members can be shown. He demonstrates the utility of the representation based on tolerance spaces by formalizing the string tied to the ratchet wheel and deriving the conclusion that the string can pull, but cannot push. To reach this conclusion, he formalizes string as a directed sequence of 'grains' such that each 'grain' transmits pulling force to the adjacent 'grain'

Perhaps, the idea proposed by Kaufmann can also be applied to the object-inside-funnel problem discussed by Davis if both examples are generalized as the problem of reasoning about effects of an influence propagating along a lengthy chain of ramifications. For a certain type of motion representable as the evolution along edges of a graph such that all transitions share some important property[14], induction can be employed to predict certain effects of motion. Similarly to proving by induction the termination of a while-loop, we do not have to consider each local transition along edges of an evolution graph separately.

[14]In the example considered by Davis, the property is that on the next step of motion, the energy diminishes; in the example considered by Kaufmann, the property is that adjacent 'grains' trasmit certain types of motion.

5 Absence of prior information

Robots may interweave planning and learning activities, but the more indefinite knowledge about environment, the more important the role of learning on the basis of sensor inputs and the less important the role of prior planning.

Different types of information can be absent. We discuss below two sub-classes.

5.1 Map construction

The first case is when there is no prior information about the environment; in particular, environment can contain unknown obstacles. Several algorithms have been proposed for learning a map from sensor data in a stationary environment: [Thrun, 1993, Buhmann *et al.*, 1995, Simmons & Koenig, 1995, Burgard *et al.*, 1996a, Burgard *et al.*, 1996b, Thrun & Buecken, 1996]. Some algorithms extract only topological information about routes in the environment, while other learn occupancy map.

[Dean *et al.*, 1995] consider map learning formulated as the problem of inferring the structure of a deterministic finite automaton from noisy observations. The states of automaton should be qualitatively different locations distinguishable by sensors. Thus, algorithm proposed in this paper construct only graph representing routes between landmarks.

[Thrun, 1993, Thrun, 1995] describe neural network approach to learning a map of an unknown two-dimensional environment. This approach is further developed in [Thrun & Buecken, 1996] where numerical occupancy grid-based map is used to build a graph similar to a Voronoi diagram [Preparata & Shamos, 1985]. This graph (built on top of the occupancy grid) represents routes between critical positions[15] in the environment. It significantly accelerates the process of navigation planning. [Thrun & Buecken, 1996] provide several arguments why combination of occupancy grid-based and topological-based approaches to navigation planning achieves better results than each of this approaches could provide alone. [Kuipers & Byun, 1991] propose hierarchical model approach towards building a map. According to their method, location-specific control algorithms are dynamically selected to control the robot's interaction with its environment. As the result of the interaction, distinctive places and paths are determined and linked to form a topological network description. The authors discuss in detail how to define distinctive places and travel paths. Once a topological network is constructed, metrical information is incrementally assimilated into local descriptions of of places and edges and finally merged into a global geometric map. Kuipers and Byun argue that their approach is more robust to systematic and random errors.

Learning a map of an unknown two-dimensional environment is a difficult problem whatever sensors are used and whatever algorithms are employed to estimate distances from robot to surrounding walls. The difficulty comes from the impossibility for sensors to look throught a wall. Consequently, measurments received from different view points should be somehow integrated into a consistent representation of an environment. [Thrun, 1993] proposes a special rule which relies on a measure of confidence for occupancy estimations. For example, if robot stays initially in a corridor, then it can estimate the space in a room as being occupied, but when it goes inside the room, it may discover that this space is empty.

[15] Critical positions are points in the map where clearance between the robot and obstacles is minimal.

[Shanahan, 1996] develops a framework which integrates learning through navigation in an initially unknown environment with reasoning about results of action sequences. He proposes to use abduction for incorporating sensor data into initially incomplete model of the world. Because his approach is logic-based, there is a way to show that his algorithm for learning an occupancy map is provably correct (in the logical sense). Shanahan emphasizes importance of dealing with noise and points out that both noisy sensors and noisy actuators can be captured using non-determinism (as an alternative to probability-based approaches). In particular, he suggests to represent the noise due to motors by considering formulas with existential quantifier over the position of the robot inside an ever expanding circle of uncertainty. It seems that this approach do simplify abductive inferences, but allows to draw only much weaker hypotheses than we would expect from statistical methods for building an occupancy map.

5.2 Skill acquisition

The second case is when there is no prior information about uncertainties in motion commands or when preliminary motion skills are unadeqate for successfull completion of a task.

There are several papers which consider how robots obtain general skills through primitive training instances or gradually learn high level descriptions from sensor inputs: [Dufay & Latombe, 1984, Zrimec & Mowforth, 1988, Morik & Rieger, 1993, Rieger, 1995].

[Dufay & Latombe, 1984] describes an approach to skill acquisition (fine motion synthesis) through inductive learning. There are two main reasons why the problem of skill acquisition occurs. First, the practical motivation is the simplification of the programming task; programming in primitive motion commands is a tedious and error prone task. Second, it is impossible to have realistic prior knowledge about the actual uncertainty of robot and its environment. The approach proposed in this paper operates in two phases: the training phase and the induction phase. During the training phase, the motion planner and the execution monitor decide on-line by interacting with the sensors which motions are to be executed. As a result, they generate multiple traces of execution with varying degree of success. Due to errors, several executions of the same plan may produce different sequences of motions. The purpose of the induction phase is to construct a program from a set of execution traces.[16] The induction phase proceeds through iterative transformation of the generated traces be rewriting rules and constructs a manipulator-level program including symbolic variables, conditional statements and loops. Authors point out that there is an analogy between their inductive task and the task of learning a grammar from training instances or synthesis of an automaton from examples. [Morik & Rieger, 1993, Rieger, 1995] address the problem of learning action-oriented perceptual features for robot navigation formulated as the problem of inferring probabilistic automata.

[Fontaine & Bidaud, 1994] describes the algorithm for finding configuration of (articulated) fingers optimal for stable grasp of objects with various shapes and grasping forces (exerted by fingers). The authors propose a way to produce a relatively simple and fast planner developed using a knowledge-based approach. To develop fine motion strategies,

[16] Each trace is recorded as a linear graph of states and motions. All the graphs have the same initial and final states.

effectors are equipped with limited decision-making capabilities. Their method uses inductive learning from the execution traces of a predefined plan, sequentially refines plans and ultimately generates a modified motion plan able to deal with uncertainties.

[Zrimec & Mowforth, 1988] describe an experiment in which a real robot is provided with a problem domain (object pushing) and learns the relationship between action and sensory information on the basis of performing controlled experiments in the world. The experiment involves random movement of a robot hand possibly interacting with an object. As a result, the object either is pushed (it changes its previous position) or remains unmoved. By using its vision system, the robot notes signals about position and orientation of the object before and after an action has taken place. All training examples were subjected to an induction algorithm which found certain regularities that relate variables in terms of cause and effect. Refining these regularities, the robot learns causal rules organized in the form of a hierarchical model: on the simplest level the robot distinguishes "push" from "no_push", on the next level it learns that "push" belongs either to "move_right" cluster, or to "move_left" cluster; finally, the robot learns more detailed specification of "push" in terms of "right_rotation", "translation" and "left_rotation".

6 Future Work

We have examined a wide range of issues relating to integration of high-level reasoning with motion planning. Because the general problem of the integration is too difficult and has many application-specific features which can be addressed only on engineering level, futher research has to focus on feasible problems rising important theoretical questions. Assuming that research will address the navigation problem for an indoor robot in an office environment, and that specifications are given to the robot in a high-level logical language (e.g., a version of GOLOG), I briefly outline a few potential (closely-connected) areas for research.

- Representation. What representation of space is useful and sufficient both for reasoning about motion and for navigation planning ? Assuming that some simple model (rectangloids, occupancy grids) does the job, on what principles declarative and procedural aspects should be divided between algorithms, that search for a path, and a logical theory, that gives high-level commands ?

- 'Intelligent' version of the `Findpath` problem: how reactive and planning aspects should be represented in the high-level language ? If robot failed to reach the destination following given task-level instructions because the selected path happens to be obstructed, how the high-level program can be modified ? If robot encounters an unexpected obstacle, but it knows that the obstacle is a movable object and robot can push it to clean the path, robot should be able to make the appropriate decision on-fly. However, the problem is not in developing of good algorithms for pushing objects, because a path can be obstructed in many different ways (for example, by people). It means that a high-level language must support local planning and reactive behaviour. Thus, appropriate enhancement of GOLOG-type languages should be developed on the basis of the appropriately extended situation calculus.

- 'Intelligent' map learning. While robot moves and its view point changes, sensor information received from surrounding walls and objects may undergo abrupt changes at certain points. This points of topological significance (corridor junctions, doors) also have properties important for navigation. It seems that map learning and sensing have to be grounded in some general theory of space with appropriate parameters (in the sense that measured values of those parameters correspond to a particular map). Given such logical theory, sonar sensor data will serve only for tuning values of parameters in the theory. Otherwise, if learning relies on the `tabula rasa` paradigm, sensor data are used both for restoring the general structure of space and for map construction. The corresponding research issues are what part of learning process should be represented in a version of the situation calculus and how ?

- Actions with probabilistic effects. Any serious research dealing with real equipment should address issues of noise and uncertainty. Because representation and reasoning about actions with probabilistic effects and about noisy sensors are issues on the reaserch frontier, their interaction with spatial reasoning should be also explored.

7 Glossary

A **blocked region** is a region inside a configuration obstacle. All configurations inside blocked region are **illegal**.

A **completely constrained** motion of a body A is a motion completely determined by a motion of other body.

A **configuration** q of a movable body $\mathcal{A}$ occupying a compact subset of $\mathbf{R}^n$ is a specification of the position and orientation of a moving Cartesian frame $\mathcal{F}_\mathcal{A}$ embedded in $\mathcal{A}$ with respect to a fixed Cartesian frame $\mathcal{F}_\mathcal{W}$ embedded in a workspace $\mathcal{W} = \mathbf{R}^n$ This definition assumes that $\mathcal{A}$ is a rigid object so that each point in $\mathcal{A}$ has fixed coordinates in $\mathcal{F}_\mathcal{A}$.

The **configuration space** of $\mathcal{A}$ is the space $\mathcal{C}$ of all possible configurations of $\mathcal{A}$. A unique arbitrarily selected configuration of $\mathcal{C}$ is called the **reference configuration** of $\mathcal{A}$. It is denoted by 0.

A **configuration obstacle** $\mathcal{CB}_i$ is a region of a **configuration space** $\mathcal{C}$ obtained by 'growing' the corresponding workspace obstacle $\mathcal{B}_i$ by the shape of the movable body $\mathcal{A}$: $\mathcal{CB}_i = \{\mathbf{q} \in \mathcal{C} \ / \ \mathcal{A}(\mathbf{q}) \cap \mathcal{B}_i \neq 0\}$.

A **configuration** space **constraint** forms boundary between a free region and a blocked region. It is defined by a pair of objects and corresponds to either a vertex or a boundary segment of one object touching the boundary of the other. The former is called vertex constraint and the latter – boundary constraint. From any configuration on the boundary of free space, points in both blocked and in free space can be reached by arbitrarily small motions.

A **free region** is a region inside the free space. All configurations inside a free region are legal. The **free space** $\mathcal{C}_{free}$ is the subset of configuration space $\mathcal{C}$ outside all configuration obstacles: $\mathcal{C}_{free} = \mathcal{C} - \bigvee_{i=1}^{k} \mathcal{CB}_i$, where k is the number of obstacles. Thus, free space is the set of all legal configurations.

The **free-flying object** assumption claims that in the absence of obstacles, any path of $\mathcal{A}$ is feasible; i.e., motions of $\mathcal{A}$ are not limited by any kinematic or dynamic constraints.

A **kinematic chain** is a chain (sequence) of interacting *kinematic pairs*. A mechanism is a kinematic chain used for the transmission of forces and motion.

A **kinematic pair** is a couple of rigid and solid objects linked together such that their relative motion is *completely constrained*.

A **legal position** of a polyhedron P (specified by a finite union of convex polyhedra p_i) is any position where each constituent polyhedron p_i intersects no obstacle and intersects no other polyhedron of P except at its specified linkage vertices. A **legal movement** of P is a continuous sequence of simultaneous translations and rotations of the polyhedra $\{p_i\}$ of P through legal positions only.

A **path** of the body $\mathcal{A}$ from the **initial** configuration $\mathbf{q}_{init}$ to the **goal** configuration $\mathbf{q}_{goal}$ is a continuous map $\tau : [0,1] \mapsto \mathcal{C}$, with $\tau(0) = \mathbf{q}_{init}$ and $\tau(1) = \mathbf{q}_{goal}$.

A **place** is a connected region of space in which all points share relevant properties. Each place is distinguished by allowable motion and contact.

A **place vocabulary** is the set of all places covering the space of interest (alternative name: **region diagram**); it can be represented by a connectivity graph. The vertices of this graph are the places, and each place is connected by edges to the places forming its boundaries, and possibly to adjacent places within the same connected region (i.e., to free-space subdivisions).

A **potential field** is a function defined on configuration space as the sum of two other functions: the attractive and repulsive potentials. The attractive potential can be defined, for example, as a quadratic well with its minimum at the goal position. The repulsive potential is a function which equals zero far away from configuration obstacles, but tends toward infinity when the distance to obstacles decreases. The negated gradient of the potential field represents an artificial force that drives robot toward the goal position while pushing it away from obstacles.

The class **PSPACE** $= \bigcup_{c \geq 1} DSPACE(n^c)$, where $DSPACE(n^c)$ denotes the class of languages accepted by a deterministic Turing machine in space n^c which is a polynomial function of n and n is the size of input words. Thus, $PSPACE$ corresponds to the problems that can be solved by a polynomial space bounded Turing machine. The following problem (quantified satisfiability) is $PSPACE$-complete: Given a boolean expression E with variables $x_1, \ldots, x_n$ decide whether $(Q_1 x_1)\ (Q_2 x_n)\ , \ldots, (Q_m x_n)\ E$ is true or not, where Q_i is either $\exists$ or $\forall$ and quantifiers can alternate arbitrarily.

A convex **polyhedral region** $\mathcal{P}$ in $\mathbf{R}^3$ is the intersection of a finite number of closed half-spaces (any plane in $\mathbf{R}^3$ decomposes the space into two half-spaces). A polyhedral region is any subset of $\mathbf{R}^3$ obtained by taking the union of a finite number of convex polyhedral regions. A **rational convex polyhedron** is specified by a finite set of linear inequalities with rational coefficients.

A **preimage** of a (sub-goal) region G is a subset P of *free space* such that if the robot's configuration is in P when it starts to follow the commanded trajectory, then it is guaranteed that (1) robot will reach G (goal reachability) and (2) it remains in G when the termination condition TC evaluates to true and motion stops (goal recognizability). It is assumed that TC has sensory inputs, the motion history and the elapsed time as its arguments and that motion stops instantaneously as soon as TC becomes true.

A **semi-algebraic set** is a subset of $\mathbf{R}^n$ whose points satisfy a *polynomial expression*. Let $P(x) \bowtie 0$ be an atomic polynomial expression over $\mathbf{R}^n$; i.e. P is a polynomial in n real variable with rational coefficients, $x \in \mathbf{R}^n$ and $\bowtie$ is any symbol in $\{=, \neq, >, <, \geq, \leq\}$. A **polynomial expression** over $\mathbf{R}^n$ is any finite boolean combination of atomic polynomial expressions over $\mathbf{R}^n$.

A **Tarski sentence** is any polynomial expression prefixed by a finite number, possibly zero, of $\exists$ and $\forall$ quantifiers applying to some of the variables in the sentence. All variables range over the set of the reals. Thus, a *polynomial expression* is a quantifier-free Tarski sentence. For example, the formula $\exists x(a \cdot x^2 + b \cdot x + c = 0)$ is a Tarski sentence if a, b, c are rational constants.

References

[10th-IJCAI, 1987] International Joint Conferences on Artificial Intelligence, Inc. 1987. *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, Milan, Italy: Morgan Kaufmann.

[14th-IJCAI, 1995] International Joint Conferences on Artificial Intelligence, Inc. 1995. *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Montrèal, Canada: Morgan Kaufmann.

[Ananthanarayanan, Goldenberg, & Mylopoulos, 1994] Ananthanarayanan, S.; Goldenberg, A.; and Mylopoulos, J. 1994. A qualitative theoretical framework for 'common-sense' based multiple contact robotics manipulation. *Robotica* 12:175–186.

[Anger & Loganantharaj, 1994] Anger, F., and Loganantharaj, R., eds. 1994. *Proceedings of the Workshop on Spatial and Temporal Reasoning at the 12th National Conference on Artificial Intelligence.*

[Anger, Guesgen, & Ligozat, 1995] Anger, F.; Guesgen, H.; and Ligozat, G., eds. 1995. *Proceedings of the Workshop on Spatial and Temporal Reasoning at the 14th International Joint Conference on Artificial Intelligence.*

[Anger, Guesgen, & van Benthem, 1993] Anger, F.; Guesgen, H.-W.; and van Benthem, J., eds. 1993. *Proceedings of the Workshop on Spatial and Temporal Reasoning at the 13th International Joint Conference on Artificial Intelligence.*

[Asher & Vieu, 1995] Asher, N., and Vieu, L. 1995. Toward a geometry of common sense – a semantics and a complete axiomatization of mereotopology. In 14th-IJCAI [1995], 846–852.

[Bäckström, 1990] Bäckström, C. 1990. Logical modeling of simplified geometrical objects and mechanical assembly processes. In Chen [1990]. 35–61.

[Barraquand & Latombe, 1991] Barraquand, J., and Latombe, J.-C. 1991. Robot motion planning: A distributed representation approach. *International Journal of Robotics Research* 10(6):629–649.

[Ben-Or, Kozen, & Reif, 1986] Ben-Or, M.; Kozen, D.; and Reif, J. 1986. The complexity of elementary algebra and geometry. *J. Comput. Syst. Sci.* 32:251–264.

[Bergacker, Williams, & Kalita, 1994] Bergacker, D.; Williams, J.; and Kalita, J. 1994. Issues in planning realistic motion from natural language instruction. In Anger and Loganantharaj [1994], 77–84.

[Brafman, Latombe, & Shoham, 1993] Brafman, R.; Latombe, J.-C.; and Shoham, Y. 1993. Towards knowledge-level analysis of motion planning.

[Brooks & Lozano-Pérez, 1985] Brooks, R., and Lozano-Pérez, T. 1985. A subdivision algorithm in configuration space for `findpath` with rotation. *IEEE Trans. Syst. Man Cybern.* 15(2):224–233.

[Brooks, 1982] Brooks, R. A. 1982. Symbolic error analysis and robot planning. *International Journal of Robotics Research* 1(4):29–68.

[Brooks, 1983] Brooks, R. A. 1983. Planning collision-free motions for pick-and-place operations. *International Journal of Robotics Research* 2(4):19–44.

[Buhmann *et al.*, 1995] Buhmann, J.; Burgard, W.; Cremers, A. B.; Fox, D.; Hofmann, T.; Schneider, F.; Strikos, J.; and Thrun, S. 1995. The mobile robot rhino. *Artificial Intelligence Magazine* 16(2):31–38.

[Burgard *et al.*, 1996a] Burgard, W.; Fox, D.; Hennig, D.; and Schmidt, T. 1996a. Estimating the absolute position of a mobile robot using position probability grids. In *Proceedings of the 13th National Conference on Artificial Intelligence*. Portland, Oregon: *AAAI*.

[Burgard *et al.*, 1996b] Burgard, W.; Fox, D.; Hennig, D.; and Schmidt, T. 1996b. Position tracking with position probability grids. Technical Report ?, Computer Science Department III, Bonn University, Germany.

[Burger & Bhanu, 1992] Burger, W., and Bhanu, B. 1992. *Qualitative Motion Understanding*, volume 184 of *The Kluwer International Series in Engineering and Computer Science*. Dordrecht: Kluwer.

[Canny & Lin, 1990] Canny, J., and Lin, M. 1990. An opportunistic global motion planner.

[Canny & Reif, 1987] Canny, J., and Reif, J. 1987. New lower bound techniques for robot motion planning problems. In *28th Annual Symposium on Foundations of Computer Science*, 49–60. Los Angeles, California: IEEE.

[Canny, Grigor'ev, & Vorobjov, 1992] Canny, J.; Grigor'ev, D.; and Vorobjov, N. 1992. Finding connected components of a semialgebraic set in subexponential time. *Applicable Algebra in Engineering, Communication and Computing* 2:217–239.

[Canny, Rege, & Reif, 1991] Canny, J.; Rege, A.; and Reif, J. 1991. An exact algorithm for kinodynamic planning in the plane. *Discrete Comput. Geom.* 6:461–484.

[Canny, 1988] Canny, J. 1988. *The Complexity of Robot Motion Planning*. ACM Doctoral Dissertation Awards, 1987. Cambridge, MA: MIT Press.

[Canny, 1993] Canny, J. 1993. Computing roadmaps of general semi-algebraic sets. *The Computer Journal* 36(5):504–514.

[Chen, 1990] Chen, S.-S., ed. 1990. *Advances in Spatial Reasoning*, volume 1 and 2. Norwood, NJ: Ablex.

[Cimatti *et al.*, 1992] Cimatti, A.; Traverso, P.; Dalbosco, S.; and Armando, A. 1992. Navigation by combining reactivity and planning. In *Proceedings of the Intelligent Vehicles Symposium*.

[Clarke, 1981] Clarke, B. 1981. A calculus of individuals based on “connection”. *Notre Dame Journal of Formal Logic* 22(3):204–218.

[Clarke, 1985] Clarke, B. 1985. Individuals and points. *Notre Dame Journal of Formal Logic* 26(1):61–75.

[Cohn *et al.*, 1994] Cohn, A.; Gotts, N.; Cui, Z.; Randell, D.; Bennett, B.; and Gooday, J. 1994. Exploiting temporal continuity in qualitative spatial calculi. In Golledge, R. G., and Egenhofer, M. J., eds., *Spatial and Temporal Reasoning in Geographical Information Systems*. Elsevier.

[Collins, 1975] Collins, G. 1975. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Automata theory and formal languages, Kaiserslautern (FRG)*, volume 33 of *Lecture Notes in Computer Science*. New York: Springer Verlag. 135–183.

[Cui, Cohn, & Randell, 1992] Cui, Z.; Cohn, A.; and Randell, D. 1992. Qualitative simulation based on a logic of space and time. In *Qualitative Reasoning about Physical Systems, 6th International Workshop, Edinburg*, 15pp.

[Cui, Cohn, & Randell, 1993] Cui, Z.; Cohn, A.; and Randell, D. 1993. Qualitative and topological relationships in spatial databases. In *Advances in Spatial Databases, 3rd Intern. Symposium on Large Spatial Databases SSD'93*, volume 692 of *Lecture Notes in Computer Science*, 296–315. Berlin: Springer.

[Davis, 1988] Davis, E. 1988. A logical framework for commonsense predictions of solid object behavior. *Int. J. of the Artificial Iintelligence in Engineering* 3(3):125–140.

[Davis, 1990] Davis, E. 1990. *Representations of commonsense knowledge*. Morgan Kaufmann.

[Davis, 1994] Davis, E. 1994. The kinematics of cutting solid objects. Technical Report tr541, Courant Institute, May 25, Department of Computer Science, New York University.

[Davis, 1995a] Davis, E. 1995a. Approximation and abstraction in solid object kinematics. Technical Report tr706, Courant Institute, September 26, Department of Computer Science, New York University.

[Davis, 1995b] Davis, E. 1995b. Approximations of shape and configuration space. Technical Report tr703, Courant Institute, September 15, Department of Computer Science, New York University.

[Dean *et al.*, 1995] Dean, T.; Angluin, D.; Basye, K.; Engelson, S.; Kaebling, L.; Kokkevis, E.; and Maron, O. 1995. Inferring finite automata with stochastic output functions and an application to map learning. *Machine Learning* 18:81–108.

[Donald & Xavier, 1990] Donald, B. R., and Xavier, P. G. 1990. Provably good approximation algorithms for optimal kinodynamic planning for Cartesian robots and open chain manipulators. In *Proc. 6th Annu. ACM Sympos. Comput. Geom.*, 290–300.

[Donald *et al.*, 1993] Donald, B.; Xavier, P.; Canny, J.; and Reif, J. 1993. Kinodynamic motion planning. *Journal of the ACM* 40(5):1048–1066. Revised version of the paper ‘On the complexity of kinodynamic planning’ published in the Proc. 29th Annu. IEEE Sympos. on Found. Comput. Sci., 1988, 306–316.

[Donald, 1987] Donald, B. 1987. A search algorithm for motion planning with six degrees of freedom. *Artif. Intell.* 31(3):295–353.

[Donald, 1988] Donald, B. 1988. A geometric approach to error detection and recovery for robot motion planning with uncertainty. *Artif. Intell.* 37(1–3):223–271.

[Donald, 1989a] Donald, B. 1989a. Planning and executing robot assembly strategies in the presence of uncertainty. Technical Report 89-1060, Cornell Computer Science Robotics Laboratory, Cornell University, Ithaca.

[Donald, 1989b] Donald, B. R. 1989b. *Error Detection and Recovery in Robotics*, volume 336 of *Lecture Notes in Computer Science*. Berlin: Springer.

[Dufay & Latombe, 1984] Dufay, B., and Latombe, J. 1984. An approach to automatic robot programming based on inductive learning. *Intern. Journal of Robotics Research* 3(4):3–20.

[Egenhofer & Al-Taha, 1992] Egenhofer, M. J., and Al-Taha, K. K. 1992. Reasoning about gradual changes of topological relationships. In Frank et al. [1992], 196–219.

[Egenhofer & J.R.Herring, 1990] Egenhofer, M., and J.R.Herring. 1990. A mathematical framework for the definition of topological relationships. In *Proc. of the 4th Intern. Symp. on Spatial Data Handling*, volume 2, 803–813.

[Egenhofer, 1989] Egenhofer, M. J. 1989. A formal definition of binary topological relationships. In Litwin, W., and Schek, H.-J., eds., *Third International Conference on Foundations of Data Organization and Algorithms*, volume 367 of *Lecture Notes in Computer Science*, 457–472. Berlin: Springer.

[Erdmann & Lozano-Pérez, 1987] Erdmann, M., and Lozano-Pérez, T. 1987. On multiple moving objects. *Algorithmica* 2:477–521.

[Erdmann, 1986] Erdmann, M. 1986. Using backprojections for fine motion planning with uncertainty. *International Journal of Robotics Research* 5(1):19–45.

[Erdmann, 1993] Erdmann, M. 1993. Randomization for robot tasks: Using dynamic programming in the space of knowledge states. *Algorithmica* 10:248–291.

[Erdmann, 1995] Erdmann, M. 1995. Understanding action and sensing by designing action-based sensors. *The International Journal of Robotics Research* 14(5):483–509. Special issue on integration among planning, sensing, and control.

[Ershov *et al.*, 1965] Ershov, Y.; Lavrov, I.; Taimanov, A.; and Taitslin, M. 1965. Elementary theories. *Russian Mathematical Surveys* 20(4):35–106. Translated from: Uspekhi Mathematicheskikh Nauk, 20, $\mathcal{N}$ 4 (124), pp. 37–108, Moscow, 1965.

[Faltings & Struss, 1992] Faltings, B., and Struss, P., eds. 1992. *Recent Advances in Qualitative Physics*. Cambridge, MA: The MIT Press.

[Faltings, Baechler, & Primus, 1989] Faltings, B.; Baechler, E.; and Primus, J. 1989. Reasoning about kinematic topology. In *IJCAI-89*, 1331–1336. Detroit, MI: International Joint Conferences on Artificial Intelligence, Inc.

[Faltings, 1987] Faltings, B. 1987. Qualitative kinematics in mechanisms. In 10th-IJCAI [1987], 436–442. Reprinted in [?].

[Faltings, 1990] Faltings, B. 1990. Qualitative kinematics in mechanisms. *Artificial Intelligence* 44:89–119.

[Faltings, 1992] Faltings, B. 1992. A symbolic approach to qualitative kinematics. *Artificial Intelligence* 56(2).

[Faltings, 1995] Faltings, B. 1995. Qualitative spatial reasoning using algebraic topology. In *International Conference on Spatial Information Theory, COSIT-95*, volume 988 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.

[Fontaine & Bidaud, 1994] Fontaine, D., and Bidaud, P. 1994. Design and implementation of an inductive learning control system for a manipulator gripper. *Robotica* 12(2):187–192.

[Forbus, Nielsen, & Faltings, 1987] Forbus, K. D.; Nielsen, P.; and Faltings, B. 1987. Qualitative kinematics: A framework. In 10th-IJCAI [1987], 430–435. Reprinted in [?].

[Forbus, Nielsen, & Faltings, 1991] Forbus, K. D.; Nielsen, P.; and Faltings, B. 1991. Qualitative spatial reasoning: the `clock` project. *Artificial Intelligence* 51(1–3):417–472.

[Forbus, 1983] Forbus, K. D. 1983. Qualitative reasoning about space and motion. In Gentner, D., and Stevens, A. L., eds., *Mental Models*. Hillsdale, N.J.: L. Erlbaum Associates. 53–73.

[Frank & Campari, 1993] Frank, A. U., and Campari, I., eds. 1993. *Spatial Information Theory. A Theoretical Basis for GIS. European Conference, COSIT'93*, volume 716 of *Lecture Notes in Computer Science*. Marciana Marina, Italy: Springer.

[Frank, Campari, & Formentini, 1992] Frank, A. U.; Campari, I.; and Formentini, U., eds. 1992. *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space. Intl. Conf. GIS—From Space to Territory*, volume 639 of *Lecture Notes in Computer Science*. Pisa: Springer.

[Fujimura, 1991] Fujimura, K. 1991. *Motion Planning in Dynamic Environments*. Computer Science Workbench. Tokio: Springer.

[Funt, 1980] Funt, B. V. 1980. Problem solving with diagrammatic representations. *Artificial Intelligence* 13(3):201–230. Reprinted in "Readings in knowledge representation", ed. Ronald J. Brachman and Hector J. Levesque, Morgan Kaufman, 1985, p.441-456.

[Gagné, Pang, & Trudel, 1995] Gagné, D.; Pang, W.; and Trudel, A. 1995. A spatio-temporal logic for 2d multi-agent problem domains. In ITESM [1995], 372–377.

[Galton, 1993] Galton, A. 1993. Towards an integrated logic of space, time, and motion. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 1550–1555. Chambéry, France: International Joint Conferences on Artificial Intelligence, Inc.

[Galton, 1995] Galton, A. 1995. Towards a qualitative theory of movement. In *International Conference on Spatial Information Theory, COSIT-95*, volume 988 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag.

[Gelsey, 1989] Gelsey, A. 1989. Automated reasoning about machine geometry and kinematics. In *The 3rd IEEE Conference on AI Applications*, 182–187. Reprinted in [?].

[Giunchiglia, Spalazzi, & Traverso, 1994] Giunchiglia, F.; Spalazzi, L.; and Traverso, P. 1994. Planning with failure. In *AIPS-94 Second International Conference on AI Planning Systems*.

[Gottschlich, Ramos, & Lyons, 1994] Gottschlich, S.; Ramos, C.; and Lyons, D. 1994. Assembly and task planning: A taxonomy. *IEEE Robotics and Automation Magazine* 1(3):4–12.

[Grigorev & Vorobjov, 1992] Grigorev, D., and Vorobjov, N. 1992. Counting connected components of a semialgebraic set. *Computational Complexity* 2:133–186.

[Grigorev, 1988] Grigorev, D. 1988. Complexity of deciding Tarski algebra. *Journal of Symbolic Computation* 5(1-2):65–108.

[Hayes, 1979] Hayes, P. J. 1979. The naive physics manifesto. In Michie, D., ed., *Expert Systems in the Micro Electronic Age*. Edinburgh University Press.

[Hayes, 1985a] Hayes, P. 1985a. Naive physics manifesto I: Ontology for liquids. In Hobbs, J., and Moore, R. C., eds., *Formal Theories of the Commonsense World.* Norwood, NJ: Ablex Publishing Corporation. 71–107.

[Hayes, 1985b] Hayes, P. 1985b. The second naive physics manifesto. In Hobbs, J. R., and Moore, R. C., eds., *Formal Theories of the Commonsense World.* Norwood, NJ: Ablex Publishing Corporation. 1–36.

[Heintz, Roy, & Solernó, 1989] Heintz, J.; Roy, M.; and Solernó, P. 1989. On the complexity of semialgebraic sets. In *Proc. of IFIP.* San Francisco: North Holland. 293–298.

[Hernández & Mukerjee, 1995] Hernández, D., and Mukerjee, A. 1995. Representation of spatial knowledge. Ijcai'95 tutorial, http://www.cs.albany.edu/ amit/tutijcai.html.

[Hernández, 1994] Hernández, D. 1994. *Qualitative Representation of Spatial Knowledge*, volume 804 of *Lecture Notes in Artificial Intelligence.* Berlin: Springer. Revised version of a 1992 Ph.D. Thesis at the Institut für Informatik, Technische Universität München.

[Herring & Egenhofer, 1995] Herring, J. R., and Egenhofer, M. J., eds. 1995. *Fourth International Symposium on Large Spatial Databases, SSD '95*, Lecture Notes in Computer Science. Springer. to appear.

[Homem de Mello & Lee, 1991] Homem de Mello, L. S., and Lee, S., eds. 1991. *Computer-aided mechanical assembly planning*, volume SECS 148 of *Kluwer international series in engineering and computer science.* Boston: Kluwer Academic Publishers.

[Hopcroft & Wilfong, 1986] Hopcroft, J., and Wilfong, G. 1986. Reducing multiple object motion planning to graph searching. *SIAM Journal on Computing* 15(3):768–785.

[Hopcroft, Schwartz, & Sharir, 1984] Hopcroft, J.; Schwartz, J.; and Sharir, M. 1984. On the complexity of motion planning for multiple independent objects: *pspace*-hardness of the 'warehouseman's problem'. *International J. of Robotics Research* 3(4):76–88.

[ITESM, 1995] ITESM. 1995. *Proceedings of the 8th International Symposium on Artificial Intelligence*, Monterrey, N.L., México: Instituto Technológico y de Estudios Superiores de Monterrey, Centro de Inteligencia Artificial.

[Joskowicz & Sacks, 1991a] Joskowicz, L., and Sacks, E. P. 1991a. Computational kinematics. *Artificial Intelligence* 51(1–3):381–416.

[Joskowicz & Sacks, 1991b] Joskowicz, L., and Sacks, E. P. 1991b. Incremental configuration space construction for mechanical analysis. In *AAAI-91*, 888–893. Menlo Park/Cambridge: American Association for Artificial Intelligence.

[Joskowicz, 1987] Joskowicz, L. 1987. Shape and function in mechanical devices. In *AAAI-87*, 611–615. Menlo Park: American Association for Artificial Intelligence.

[Kambhampati & Davis, 1986] Kambhampati, S., and Davis, L. 1986. Multiresolution path planning for mobile robots. *IEEE Journal of Robotics and Automation* 2:135–145.

[Kaufman, 1991] Kaufman, S. 1991. A formal theory of spatial reasoning. In *Proceedings of the 2nd conference on Principles of Knowledge Representation and Reasoning*, 347–356. San Mateo, CA: Morgan Kaufmann.

[Khatib, Craig, & Lozano-Pérez, 1992] Khatib, O.; Craig, J.; and Lozano-Pérez, T. 1992. *The Robotics Review*, volume 2. Cambridge, Mass.: MIT Press. Vol.1 was published in 1989.

[Khatib, 1986] Khatib, O. 1986. Real-time obstacle avoidance for manipulators and mobile robots. *Internat. J. of Robot. Research* 5(1):90–98.

[Koga *et al.*, 1994] Koga, Y.; Kondo, K.; Kuffner, J.; and Latombe, J. 1994. Planning motions with intentions. In *Proceedings of SIGGRAPH'94, Computer Graphics Proceedings*, 395–408.

[Kramer, 1992] Kramer, G. A. 1992. *Solving Geometric Constraint Systems: A Case Study in Kinematics.* Cambridge, Massachusetts: The MIT Press.

[Kuipers & Byun, 1991] Kuipers, B. J., and Byun, Y.-T. 1991. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems* 8(1-2):47–63.

[Kuipers & Levitt, 1988] Kuipers, B. J., and Levitt, T. 1988. Navigation and mapping in large-scale space. *Artificial Intelligence Magazine* 9(2):25–43. Reprinted in [Chen, 1990], vol.2, pp.207–251.

[Kuipers, 1990] Kuipers, B. J. 1990. Three chapters on the tour model. In Chen [1990]. 169–251.

[Kusiak, 1990] Kusiak, A. 1990. *Intelligent manufacturing systems.* Englewood Cliffs, N.J: Prentice-Hall.

[Latombe, Lazanas, & Shekhar, 1991] Latombe, J.-C.; Lazanas, A.; and Shekhar, S. 1991. Robot motion planning with uncertainty in control and sensing. *Artificial Intelligence* 52(1):1–47.

[Latombe, 1991] Latombe, J.-C. 1991. *Robot Motion Planning*, volume 124 of *The Kluwer International Series in Engineering and Computer Science.* Boston: Kluwer Academic Publishers.

[Lazanas, 1994] Lazanas, A. 1994. *Reasoning about Uncertainty in Robot Motion Planning.* Ph.D. Dissertation, Dep. of Computer Sceince of Stanford University, Stanford, California.

[Lesperance *et al.*, 1994] Lesperance, Y.; Levesque, H.; Lin, F.; Marcu, D.; Reiter, R.; and Scherl, R. 1994. A logical approach to high-level robot programming – a progress report. In Kuipers, B., ed., *Control of the Physical World by Intelligent Systems*, 79–85.

[Lozano-Pérez & Winston, 1977] Lozano-Pérez, T., and Winston, P. 1977. LAMA: A language for automatic mechanical assembly. In *5th Intern. Joint Conf. on AI*, 710–716.

[Lozano-Pérez *et al.*, 1992] Lozano-Pérez, T.; Jones, J. L.; Mazer, E.; and O'Donnell, P. A. 1992. *HANDEY: A Robot Task Planner.* Cambridge (MA, USA): The MIT Press.

[Lozano-Pérez, Mason, & Taylor, 1984] Lozano-Pérez, T.; Mason, M. T.; and Taylor, R. 1984. Automatic synthesis of fine-motion strategies for robots. *Internat. J. Robot. Res.* 3(1):3–24.

[Lozano-Pérez, 1981] Lozano-Pérez, T. 1981. Automatic planning of manipulator transfer movements. *IEEE Trans. Syst. Man Cybern.* SMC-11(10):681–698.

[Lozano-Pérez, 1982] Lozano-Pérez, T. 1982. Task planning. In Brady, M.; J.M.Hollerbach; T.L.Johnson; Lozano-Pérez, T.; and Mason, M., eds., *Robot Motion.* Cambridge, Massachusetts: MIT Press. 473–498.

[Lozano-Pérez, 1983] Lozano-Pérez, T. 1983. Spatial planning: A configuration space approach. *IEEE Transactions on Computers* C-32(2):108–120.

[Mason, 1981] Mason, M. 1981. Compliance and force control for computer controlled manipulators. *IEEE Transactions in System Man and Cybernetics* 11(6):418–432.

[Morik & Rieger, 1993] Morik, K., and Rieger, A. 1993. Learning action-oriented perceptual features for robot navigation. Report 3, University of Dortmund, Lehrstuhl Informatik VIII - AI.

[Mukerjee, Manish, & Praveen, 1995] Mukerjee, A.; Manish, A.; and Praveen, B. 1995. A qualitative discretization for 3d contact motions. In 14th-IJCAI [1995], 915–921.

[Myers & Wilkins, 1995] Myers, K., and Wilkins, D. 1995. Reasoning about locations in theory and practice. Technical report, Artificial Intelligence Center, SRI International.

[Nielsen, 1988] Nielsen, P. E. 1988. A qualitative approach to mechanical constraint. In *AAAI-88*, 270–274. Menlo Park: American Association for Artificial Intelligence. Reprinted in [?].

[Nilsson, 1969] Nilsson, N. J. 1969. A mobile automaton: An application of artificial intelligence techniques. In *Proc. of the Intern. Joint Conf. on Artif. Intelligence*, 509–520.

[Nilsson, 1980] Nilsson, N. J. 1980. *Principles of artificial intelligence.* Los Altos: Morgan Kaufmann.

[Olivier, Ormsby, & Keiichi, 1995] Olivier, P.; Ormsby, A.; and Keiichi, N. 1995. Higher-pair kinematics using occupancy arrays. In Anger et al. [1995], 65–74.

[Pang & Trudel, 1994] Pang, W., and Trudel, A. 1994. A point based spatio-temporal logic. In Anger and Logan antharaj [1994], 145–150.

[Pang & Trudel, 1995] Pang, W., and Trudel, A. 1995. A logic for representation and reasoning about robotic navigation on a shop floor. In ITESM [1995], 139–146.

[Pianesi & Varzi, 1994] Pianesi, F., and Varzi, A. 1994. Mereotopological construction of time from events. In Cohn, A., ed., *Proceedings of 11th Europen Conference on Artificial Intelligence*, 396–400. John Wiley & Sons, Ltd.

[Popplestone, Liu, & Weiss, 1990] Popplestone, R. J.; Liu, Y.; and Weiss, R. 1990. Group theoretic approach to assembly planing. *AI Magazine* 11(1):82–97.

[Preparata & Shamos, 1985] Preparata, F. P., and Shamos, M. I. 1985. *Computational geometry: an introduction.* Berlin: Springer-Verlag.

[Qin, Cameron, & McLean, 1995] Qin, C.; Cameron, S.; and McLean, A. 1995. Towards efficient motion planning for manipulators with complex geometry. In *Proc. of the IEEE Intern. Symposium on Assembly and Task Planning*, 207–212.

[Rajagopalan, 1994a] Rajagopalan, R. 1994a. A model for integrated qualitative spatial and dynamic reasoning in physical systems. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1411–1417. Cambridge: American Association for Artificial Intelligence.

[Rajagopalan, 1994b] Rajagopalan, R. 1994b. On shape abstractions for qualitative spatial reasoning. In Anger and Logan antharaj [1994], 151–155.

[Randell & Cohn, 1992] Randell, D., and Cohn, A. 1992. Naive topology: Modelling the force pump. In B.Faltings and P.Struss [1992]. 177–192.

[Reif & Sharir, 1994] Reif, J., and Sharir, M. 1994. Motion planning in the presence of moving obstacles. *Journal of the ACM* 41(4):764–790. Published initially in the Proc. of 25th FOCS'79 IEEE Symposium.

[Reif & Storer, 1994] Reif, J. H., and Storer, J. A. 1994. A single-exponential upper bounds for finding shortest paths in three dimensions. *Journal of the ACM* 41(5):1013–1019.

[Reif, 1987] Reif, J. H. 1987. Complexity of the generalized movers problem. In Hopcroft, J.; Schwartz, J.; and Sharir, M., eds., *Planning, Geometry and Complexity of Robot Motion.* Norwood, NJ: Ablex Pub. Corp. 267–281. Published initially in the Proc. of 20th FOCS'79 IEEE Symposium, p.421-427.

[Renegar, 1992] Renegar, J. 1992. On the computational complexity and geometry of the first-order theory of the reals. *Journal of Symbolic Computation* 13:255–352.

[Rieger, 1995] Rieger, A. 1995. Inferring probabilistic automata from sensor data for robot navigation. Report 18, University of Dortmund, Lehrstuhl Informatik VIII - AI.

[Rimon & Koditschek, 1992] Rimon, E., and Koditschek, D. 1992. Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation* 8(5):501–518.

[Rogers, 1956] Rogers, H. 1956. Certain logical reductions and decision problems. *Annals of Mathematics* 64(2):264–284.

[Roos & Noltemeier, 1991] Roos, T., and Noltemeier, H. 1991. Dynamic Voronoi diagrams in motion planning. In *Computational Geometry — Methods, Algorithms and Applications: Proc. Internat. Workshop Comput. Geom. CG '91*, volume 553 of *Lecture Notes in Computer Science.* Springer-Verlag. 227–236.

[Samet, 1990a] Samet, H. 1990a. *Applications of spatial data structures: computer graphics, image processing, and GIS.* Reading, Mass.: Addison-Wesley.

[Samet, 1990b] Samet, H. 1990b. *The design and analysis of spatial data structures.* Addison-Wesley series in computer science. Reading, Mass.: Addison-Wesley.

[Schwartz & Sharir, 1983] Schwartz, J. T., and Sharir, M. 1983. On the piano movers' problem, ii: General techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics* 4:298–351. Reprinted in [Schwartz, Sharir, & Hopcroft, 1987].

[Schwartz & Yap, 1987] Schwartz, J., and Yap, C.-K., eds. 1987. *Algorithmic and geometric aspects of robotics*, volume 1 of *Advances in robotics.* Hillsdale (NJ): Lawrence Erlbaum.

[Schwartz, Sharir, & Hopcroft, 1987] Schwartz, J. T.; Sharir, M.; and Hopcroft, J., eds. 1987. *Planning, geometry, and complexity of robot motion.* Ablex series in artificial intelligence. Ablex.

[Shaffer, 1992] Shaffer, C. 1992. A real-time robot arm collision avoidance system. *IEEE Journal of Robotics and Automation* 8(2):149–160.

[Shanahan, 1995] Shanahan, M. 1995. Default reasoning about spatial occupancy. *Artificial Intelligence* 74(1):147–163.

[Shanahan, 1996] Shanahan, M. 1996. Robotics and the commonsense informatic situation. In *Working Notes for the 3rd Symposium on Formalization of Commonsense Reasoning*, 186–198.

[Simmons & Koenig, 1995] Simmons, R., and Koenig, S. 1995. Probabilistic robot navigation in partially observable environments. In 14th-IJCAI [1995], 1080–1087.

[Thrun & Buecken, 1996] Thrun, S., and Buecken, A. 1996. Integrating grid-based and topological maps for mobile robot navigation. Technical report, University of Bonn.

[Thrun, 1993] Thrun, S. 1993. Exploration and model building in mobile robot domain. In *Proc. of the IEEE Intern. Conf. on Neural Networks.* San Francisco, CA: IEEE Press.

[Thrun, 1995] Thrun, S. 1995. An approach to learning mobile robot navigation. *Robotics and Autonomous Systems* 15:301–319.

[Topaloglou, 1991] Topaloglou, T. 1991. Representation and management issues for large spatial knowledge bases. Depth Oral paper submitted in partial fulfillment of the requirements for the PhD degree.

[Traverso *et al.*, 1994] Traverso, P.; Cimatti, A.; Spalazzi, L.; Giunchiglia, E.; and Armando, A. 1994. Mrg: Building planners for real world complex applications. *Applied Artificial Intelligence* 8(3):333–357.

[Udupa, 1977] Udupa, S. M. 1977. Collision detection and avoidance in computer controlled manipulators. In *5th Intern. Joint Conf. on Artif. Intell.*, 737–748.

[Vieu, 1993] Vieu, L. 1993. A logical framework for reasoning about space. In Frank and Campari [1993], 25–35.

[Waldinger, 1977] Waldinger, R. 1977. Achieving several goals simultaneously. In Elcock, E., and Michie, D., eds., *Machine Intelligence, vol. 8: Machine Representations of Knowledge*. Chichester, UK: Ellis Horwood. 94–136.

[Weld & de Kleer, 1990] Weld, D. S., and de Kleer, J., eds. 1990. *Readings in Qualitative Reasoning about Physical Systems*. San Mateo, CA: Morgan Kaufmann.

[Wilfong, 1988] Wilfong, G. 1988. Motion planning in the presence of movable obstacles. In *Proc. of the 4th ACM Sympos. on Computational Geometry*, 279–288.

[Wilson & Latombe, 1994] Wilson, R. H., and Latombe, J.-C. 1994. Geometric reasoning about mechanical assembly. *Artificial Intelligence* 71:371–396.

[Zrimec & Mowforth, 1988] Zrimec, T., and Mowforth, P. 1988. An experiment in generating deep knowledge for robots. In Gero, J., ed., *Artificial Intelligence in Engineering: Robotics and Processes*. Amsterdam: Elsevier. 21–33.