2-19-1991

Toward a Science of Robot Planning and Control

Daniel E. Koditschek
University of Pennsylvania, kod@seas.upenn.edu

Recommended Citation

NOTE: At the time of publication, author Daniel Koditschek was affiliated with Yale University. Currently, he is a faculty member in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

Toward a Science of Robot Planning and Control

Abstract
Programming machines to operate flexibly and autonomously in the physical world seems to require a sophisticated representation that encodes simultaneously the nature of a task, the nature of the environment within which the task is to be performed, and the nature of the robot's capabilities with respect to both. We seek a scientific methodology of robot task encoding that encompasses the desired behavioral goals and environmental conditions as well. The methodology must balance the need for flexible expression of abstract human goals against the necessity of a eliciting a predictable response from the commanded machine. This talk focuses on the problem of motion planning as an example of how we propose to say what we mean to a robot and to know what we have said.

For more information: Kod*Lab

Disciplines
Electrical and Computer Engineering | Engineering | Systems Engineering

Comments
NOTE: At the time of publication, author Daniel Koditschek was affiliated with Yale University. Currently, he is a faculty member in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

Toward a Science of Robot Planning and Control

Extended Abstract of a Talk Delivered at the AAAS Annual Meeting
Technical Workshop on Robotics and Mathematics
Washington DC, February 14-19, 1991

Daniel E. Koditschek *

Center for Systems Science
Yale University, Department of Electrical Engineering

1 Introduction

Programming machines to operate flexibly and autonomously in the physical world seems to require a sophisticated representation that encodes simultaneously the nature of a task, the nature of the environment within which the task is to be performed, and the nature of the robot's capabilities with respect to both. We seek a scientific methodology of robot task encoding that encompasses the desired behavioral goals and and environmental conditions as well. The methodology must balance the need for flexible expression of abstract human goals against the necessity of a eliciting a predictable response from the commanded machine. This talk focuses on the problem of motion planning as an example of how we propose to say what we mean to a robot and to know what we have said.

2 Task Encoding

Roughly speaking, control theory has traditionally postulated a plant and a reference trajectory, and addressed the problem of how to force the plant to "track" the desired trajectory. Speaking even more roughly, artificial intelligence has traditionally postulated some abstract (generally human) goal and a symbol processing machine, and addressed the question of how to represent the goal within that symbolic system, sometimes considering the problem of realizing that goal symbolically as well. There is an obvious gulf between the (traditionally static) symbol system and (the usually physical) dynamical plant: roughly speaking, one must ask "where did the reference trajectory come from?"

2.1 Desiderata

Arguably, then, the business of "intelligent control theory" should concern the generation of goal representations that admit an automatic synthesis of controllers whose effect upon the plant is provably correct. Thus, the term "task encoding" is intended to encompass

1. a programming language based upon the dynamical geometry of the plant and its environment;
2. a synthesis procedure capable of generating controllers for the plant automatically from the programming language;
3. a proof that the plant together with the controller effects the correct behavior within its environment;
4. a generalization scheme for re-using, in "equivalent" situations, task specifications and/or the controllers that result.

The work my students and I have completed in the field of robotics represents a very tentative and incomplete effort toward a theory of task encoding for a greatly restricted class of plants and environments. We have as yet no formal language for any setting. In some areas our correctness proofs lag far behind our experimental observations, in others our proofs require such unrealistic assumptions that no fruitful experiments have yet been attempted. Our notions of generalization remain narrowly specialized. Nevertheless, the program of task encoding, however ambitious or even, possibly, wrong-headed, seems to represent one of the few clearly articulated statements of what it might mean to have a scientific program of research in "intelligent control."

2.2 Representation

One of the reasons that the program outlined above is difficult to put into practice is the general absence of models that are sufficiently accurate to represent the phenomena of interest but sufficiently simple to admit some hope of subsequent analysis.

Of the fundamental representation issues introduced at the beginning of this paper, only the last may be said to be scientifically well founded at

*This work was supported in part by the National Science Foundation under a Presidential Young Investigator Award.

present: the robot may be effectively modelled as a nonlinear control system using techniques of 19th century physics. Accepting the standard physical model of the robot as a controlled dynamical system, any language for prescribing its behavior must ultimately be resolved into force commands at the actuator level. The resulting behavior is then governed by certain nonlinear differential equations and it makes sense to model task and environment in terms amenable to analysis by dynamical systems theory.

The matter of representation now becomes knottier. In all but the most trivial instances, the robot's desired behavior involves interaction with an environment, $\mathcal{E}$, that must itself possess geometric and dynamical properties. Moreover, in the context of particular tasks, various aspects of the robot's operation in the environment will give rise to a new set, $\mathcal{P}$, that might be called the "planning set", within which particular goals may be formally represented. Finally, a robot operating in a specified environment might be assigned a variety of tasks. The specific task desired — an abstraction meaningful initially only to its human originator — must be encoded in terms that relate to the robot in its environment. Thus, within the context of the planning set there must be devised a formal representation of the desired behavior — the "encoding". It is not at all clear how to tell a robot to "fold the laundry" or "scramble the eggs" or "make the bed". For such tasks neither the environment, $\mathcal{E}$, nor the appropriate planning space, $\mathcal{P}$, nor the task encoding seem very obvious. In contrast, this talk presents a program of task encoding over a domain wherein the ordinarily confusing issues of representation seem relatively straightforward.

3 The Navigation Problem

Let a robot move in a cluttered but perfectly known workplace. There is a particular location of interest and it is desired that the robot move to this location from anywhere else in the workplace without colliding with the obstacles present.

3.1 Representation

The constituent pieces of the problem seem readily apparent in this case. The robot model has already been introduced. The environment, $\mathcal{E}$, is simply the workplace — a subset of Euclidean 3-space remaining after the obstacles are removed. Contained within the robot's configuration space is the *free space*, $\mathcal{F}$ — the set of all robot placements which do not involve intersection with any of the "obstacles" cluttering the workplace. The appropriate planning set, $\mathcal{P}$, for this problem is now clear: it is the phase space formed over $\mathcal{F}$, that is, the union of all the robot's configuration space velocity vectors taken over each configuration in $\mathcal{F}$. For present purposes this may be modelled as a smooth manifold with boundary (but see [20] for the case of sharp corners). The task also seems straightforward to represent: a particular navigation problem results from the choice of one particular destination point in the interior of the freespace. The goal set, $\mathcal{G}$, is a singleton: the destination point at zero velocity. The problem is now to find a feedback controller under whose influence the robot's state will approach $\mathcal{G}$ from as large as set of initial configurations as possible while remaining in $\mathcal{P}$.

A few caveats are in order before proceeding. First, it is entirely likely that the robot's freespace is not connected — that is, there may be no collision free path from some legal configurations to the destination. In the more traditional version of this problem, the navigation problem includes the decision task of whether a particular initial configuration is in fact included in the same connected component of $\mathcal{F}$ as the destination. In the present formulation the robot must arrive (with probability one) at the goal if a path exists. Thus one can only conclude (with probability one) that no such path exists for a particular initial configuration only after the robot's motion under the controller ceases at some spurious location. Second, a constructive representation of the planning space, $\mathcal{P}$, may be very difficult to obtain in practice, even when $\mathcal{E}$ is perfectly known (which, of course, it might not be in the real world). Yet this work presumes that exact information concering the boundary components of $\mathcal{P}$ is available.

3.2 Navigation Functions

Motivated by Lord Kelvin's assurance that dissipative mechanical systems end up at the local minima of the potential field, a great deal of interest in robotics has centered around the construction of artificial potential fields to encode navigation problems. Initiated by Khatib a decade ago [8], the idea of using artificial potential functions for robot task description and control was adopted or re-introduced independently by a number of researchers [17, 1, 18]. Since the interest in artificial potential functions originally emerged within the robotic control community, it is perhaps not surprising that little attention was paid to the algorithmic issues of global path planning in this literature. The question of whether the method could be used to guarantee the construction of a path between any two points in a path-connected space remained unexplored. Yet it is exactly this kind of global property that would lend autonomy from "higher level" intelligence to the controller.

A Practicable Global Stability Mechanism In the present context, the utility of artificial potential functions for path planning rests upon the possibility of deducing global stability properties from local computations. Because the potential function serves as a global Lyapunov function for its gradient vector field, it is easy to see that the minima of a gradient system (that satisfies certain regularity conditions) will attract almost all trajectories [7, 12]. Of course, the condition for a minimum is a local one that may be constructively checked via calculus and algebraic computation. Thus, if it can be assured that there is only one minimum and that it coincides with the desired destination then a potential function serves as a global path planner on the freespace, $\mathcal{F}$. Of course, the appropriate planning space is $\mathcal{P}$, the space of legal configurations and all their possible velocities. But a slight extension to Lord Kelvin's century old results on energy dissipation suffices to make the same machinery work with a suitably designed controller for the robot on $\mathcal{P}$ [12].

Existence Gradually, there seems to have emerged a common awareness of several fundamental problems with the potential function methodology First, researchers inevitably discovered through simulations or actual implementations that progressive summation of additional obstacles often lead to spurious minima and their accompanying local basins of attraction into which the robot would generally "stall out" long before acheiving the desired destination. Second, the infinite value of the artificial potentials required to prevent trajectories of the ultimate mechanically controlled system from crashing through obstacle boundaries obviously could not be achieved in the physical world and there were no clear guarantees as to when the saturation torque levels of the robot's actuators would indeed suffice to prevent collisions. Thus, an artificial potential functon need satisfy a list of technical conditions in order to give rise to a bounded torque feedback controller that guarantees convergence to the goal state, $\mathcal{G}$, from almost every initial configuration. This list comprises comprises the notion of a *navigation function* introduced to the literature two years ago [19].

The question immediately arises whether such desirable features may be achieved in general. In fact, the answer is affirmative: smooth navigation functions exist on any compact connected smooth manifold with boundary [16]. Thus, in any problem involving motion of a mechanical system through a cluttered space (with perfect information and no requirement of physical contact) if the problem may be solved at all, we are guaranteed that it may be solved by a navigation function. There remains the engineering problem of how to construct such functions

Invariance The importance of coordinate changes and their invariants is by now a well known theme in control theory. Roughly speaking, these notions formalize the manner in which two apparently different problems are actually the same. Their most familiar instance is undoubtedly encountered in the category of linear maps on linear vector spaces whose invariants (under changes of basis) determine closed loop stability. Of course, many other instances may be found in the control literature and, more recently, the utility of coordinate changes in robotics applications has been proposed independently by Brockett [2] as well.

The relevant invariant in navigation problems is the topology of the underlying configuration space [10]. In this regard, the significant virtue of the navigation function is that its desirable properties are invariant under diffeomorphism [16]. Thus, instead of building a navigation function for each particular problem, we are encouraged to devise "model problems", construct the appropriate model navigation functions, and then "deform" them into the particular details of a specified problem.

4 The Construction of Navigation Functions

4.1 A "Model" Problem

A "Euclidean sphere world" is a compact connected subset of $\mathbf{E}^n$ whose boundary is the disjoint union of a finite number, say $M+1$, of $(n-1)$-spheres. We suppose that perfect information about this space has been furnished in the form of $M+1$ center points $\{q_i\}_{i=0}^M$ and radii $\{\rho_i\}_{i=0}^M$ for each of the bounding spheres. There are two new ideas in our artificial potential function construction. First, we avoid spurious minima by multiplying the constituent functions together rather than summing them up. Namely, the "bad" set of obstacle boundaries to be avoided is encoded by the product function, $\beta : \mathcal{M} \to [0, \infty)$ is,

$$\beta \triangleq \Pi_{i=0}^M \beta_i,$$

where

$$\beta_0 \triangleq \rho_0^2 - \|q\|^2 \quad ; \quad \beta_j \triangleq \|q-q_j\|^2 - \rho_j^2 \quad j = 1 \dots M$$

are the outer boundary and inner obstacle functions, respectively. The good set, the desired destination, q_d is represented by an ordinary Hook's Law potential, $\gamma \triangleq \|q-q_d\|^{2k}$, raised to an even power and the rough syntax "go to $\gamma = 0$ and do not go to $\beta = 0$" is encoded by the intuitively obvious product

$$\hat{\varphi} \triangleq \frac{\gamma}{\beta}.$$

Of course, $\hat{\varphi}$ is unacceptable since it is unbounded. The second new idea at work is to produce a bounded potential and gradient by a smooth "squashing" function,

$$\sigma(x) \triangleq \frac{x}{1+x}.$$

Note that the composition

$$\sigma \circ \hat{\varphi} = \frac{\gamma}{\gamma + \beta}$$

is everywhere smooth and bounded, and attains its maximal height of unity only on the boundary components of the configuration space. For technical reasons we find it necessary to take the k^{th} root of this ratio with the following result.

Theorem 1 ([16]) *If the configuration space, $\mathcal{J}$, is a Euclidean sphere world then for any finite number of obstacles, and for any destination point in the interior of $\mathcal{J}$,*

$$\varphi = \sigma_d \circ \sigma \circ \hat{\varphi} = \left(\frac{\gamma^k}{\gamma^k + \beta} \right)^{\frac{1}{k}}, \qquad (1)$$

has no degenerate critical points and attains the its maximal value of unity on the boundary, $\partial \mathcal{J}$. Moreover, there exists a positive integer N such that for every $k \geq N$, φ has one and only one minimum on $\mathcal{J}$.

The function, N, on which the theorem depends is given explicitly in [16].

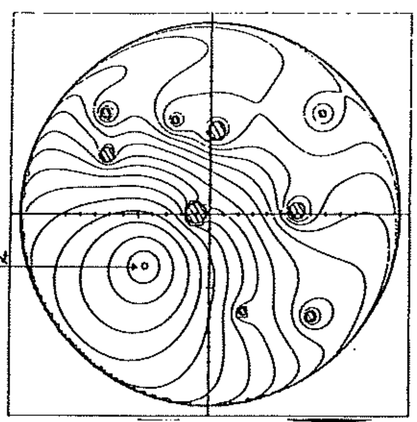

Figure 1: Planar sphere world with nine internal obstacles [23]. Contour lines denote the level curves of a navigation function constructed according to Theorem 1.

4.2 A Class of Coordinate Transformations

A *star shaped set* is a diffeomorph of a Euclidean n-disk, $\mathcal{D}^n$ possessed of a distinguished interior *center point* from which all rays intersect its boundary in a unique point. A *star world* is a compact connected subset of $\mathbf{E}^n$ whose boundary is the disjoint union of a finite number of star shaped set boundaries. Now suppose the availability of an implicit representation for each boundary component: that is, let β_i be a smooth scalar valued function that is positive outside, negative inside, and vanishes on the boundary of the i^{th} obstacle. Assume, moreover, that a known center point location, q_j has been specified for each obstacle as well. Further geometric information required in the construction to follow is detailed in the chief reference for this work [21]. A suitable Euclidean sphere world model, $\mathcal{M}$, is explicitly constructed from this data. That is, one determines (p_j, ρ_j), the center and radius of a model j^{th} sphere, according to the center and minimum "radius" (the minimal distance from q_j to the j^{th} obstacle) of the j^{th} star shaped obstacle.

A transformation, $h : \mathcal{M} \to \mathcal{F}$, may now be constructed in terms of the given star world and the derived model sphere world geometrical parameters as follows. Denote the "j^{th} omitted product", $\Pi_{j=0}^{M} \beta_j$ as $\bar{\beta}_j$. The "j^{th} analytic switch", $\sigma_j \in C^{\omega}[\mathcal{F}, \mathbb{R}]$,

$$\sigma_j(q, \lambda) \triangleq \frac{x}{x + \lambda} \circ \frac{\gamma_d \bar{\beta}_j}{\beta_j} = \frac{\gamma_d \bar{\beta}_j b}{\gamma_d \bar{\beta}_j b + \lambda \beta_j},$$

(where λ is a positive constant) attains the value one on the j^{th} boundary and the value zero on every other boundary component of $\mathcal{F}$. The "j^{th} star set deforming factor", $\nu_j \in C^{\omega}[\mathcal{F}, \mathbb{R}]$,

$$\nu_j(q) \triangleq \rho_j \frac{1 + \beta_j(q)}{\|q - q_j\|},$$

scales the ray starting at the center point of the j^{th} obstacle, q_j, through its unique intersection with that obstacle's boundary in such a way that q is mapped to the corresponding point on the j^{th} model obstacle — a suitable sphere. The overall effect is that the complicated star shaped obstacle is is "deformed along the rays" originating at its center point onto the corresponding sphere in model space.

The *star world transformation* is now given as

$$\begin{aligned} h_\lambda(q) \ &\triangleq \textstyle\sum_{j=0}^{M} \sigma_j(q, \lambda) \left[\nu_j(q) \cdot (q - q_j) + p_j \right] \\ &\quad + \sigma_d(q, \lambda) \left[(q - q_d) + p_d \right], \end{aligned} \qquad (2)$$

where σ_j is the j^{th} analytic switch, σ_d is defined by

$$\sigma_d \triangleq 1 - \sum_{j=0}^{M} \sigma_j, \qquad (3)$$

and ν_j is the j^{th} star set deforming factor. The "switches", make h look like the j^{th} deforming factor in the vicinity of the j^{th} obstacle, and like the identity map away from all the obstacle boundaries. With some further geometric computation we are able to prove the following.

Theorem 2 ([21]) *For any valid star world, $\mathcal{F}$, there exists a suitable model sphere world $\mathcal{M}$, and a positive constant Λ, such that if $\lambda \geq \Lambda$, then*

$$h_\lambda : \mathcal{F} \rightarrow \mathcal{M},$$

is an analytic diffeomorphism.

Thus, if φ is a navigation function on $\mathcal{M}$, the construction of h_λ automatically induces a navigation function on $\mathcal{F}$ via composition, $\tilde{\varphi} \stackrel{\Delta}{=} \varphi \circ h_\lambda$.

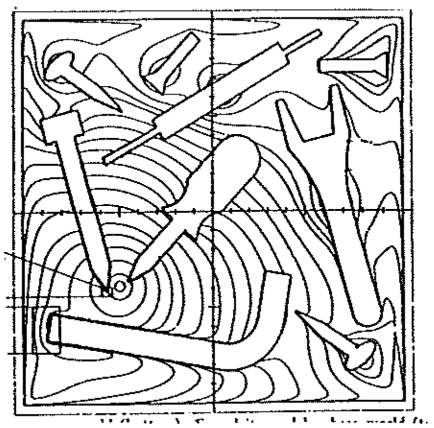

Figure 2: Planar star world with nine internal obstacles [23]. The contour lines are level curves of a navigation function induced by diffeomorphism according to Theorem 2, modified to take account of the "sharp corners" [22]. The model sphere world is depicted in the previous Figure 1.

4.3 Navigation Functions for Geometrically Complicated Spaces

In a recent paper [20], we show how to extend significantly this class of coordinate transformations. Briefly, consider an obstacle $\mathcal{O}_i$ which is a union of several intersecting stars. The arrangement of the stars in $\mathcal{O}_i$ can be partially described by a graph. If the graph is a tree, and the geometric arrangement of the daughters to the parent stars in the tree satisfies certain other regularity assumptions [20] then say that the obstacle is a *tree of stars;* a *forest of trees of stars* is a freespace, $\mathcal{F}$, consisting of the disjoint union of a finite number of trees of stars. It can be shown that any deformed sphere world can be approximated arbitrarily closely by a suitable forest of trees of stars [23]. Such a forest, $\mathcal{F}$ has a *purged version* , $\hat{\mathcal{F}}$, defined to be $\mathcal{F}$ with the leaves in trees consisting of more than one star filled-in and "reattached" to $\mathcal{F}$. Using the ideas presented above, we have shown how to define a change of coordinates from any forest, $\mathcal{F}$, to its purged version, $\hat{\mathcal{F}}$, [20]. Successive purged versions of a forest result eventually in a star world. Thus by composing successively such "purging transformations", we change coordinates from the original forest of trees of stars to a star world on which a navigation function can be constructed as described above.

Figure 3 depicts a two-dimensional (the results, of course, work in arbitrary dimensions) forest of stars resembling a building floor plan. There are three internal tree-like obstacles, and the depth of the deepest tree is $d = 4$. According to the method described above, the purging transformation, f_{Λ_i}, is applied d times, until a space whose obstacles are the roots of the original trees is obtained. This space is a star world: the the previously constructed star-world to sphere-world transformation [21] may now be used to to obtain the corresponding model sphere world, $\mathcal{M}$, on which the simple navigation function may be used.

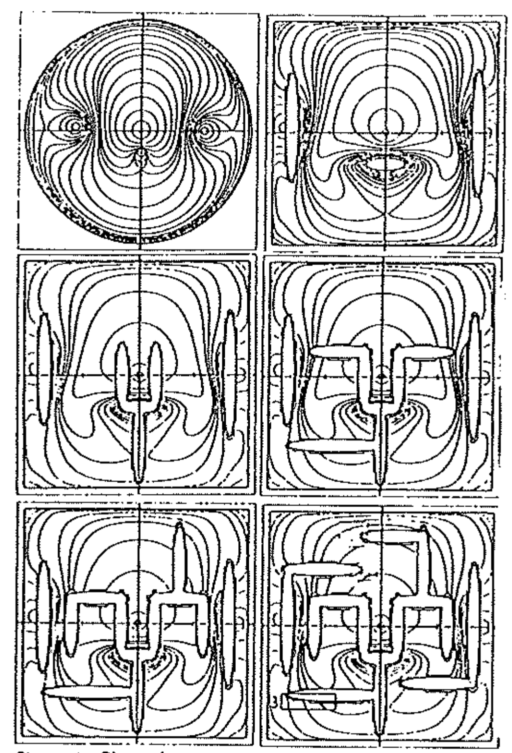

Figure 3: Planar forest of stars with three internal tree-like obstacles (bottom right), its "purged" versions, and its model sphere world (top left). [20]

5 Conclusion

This paper has presented an example of the task encoding program that underlies all of our intelligent controls work within the Yale Robotics Laboratory. Its focus, the navigation function represents the broadest scope of application we have yet found for any particular encoding. From the point of view of traditional control theory this merely represents an extension of proportional-derivative feedback techniques to settings wherein global convergence cannot obtain from linear methods because the state space is topologically distinct from a Euclidean vector space. Nevertheless, this simple notion seems to have considerable power in applications [11] beyond the a control-theoretic reformulation of the famous "piano movers' problem," [24] treated here.

In the next few years it will be essential to develop an adaptive version of these techniques in order to tolerate uncertainty in à priori knowledge of the environment. Moreover, some means of classifying the topological equivalence class of reasonably common free configuration spaces will be essential if the constructive nature of the approach is to be generalized further. Most broadly, the chief need is to remove enough à priori assumptions in the theory that reasonable physical experiments may begin.

This program of research has been extended beyond the purely geometric task domain considered here. We have succeeded in encoding certain primitive dynamically dexterous tasks as throwing, catching, and juggling using a geometric formalism [6, 5], that seems to admit correctness proofs in simple cases [3, 4]. The stability mechanism underlying the success of these computational structures has been understood only recently within the "chaos" literature, and seems to generalize to successful hopping algorithms as well [9].

There are present plans to extend the purview of this methodology still further. Planning and control of automated assembly may be cast as the problem of controlling a nonholonomically constrained mechanical system [13, 14]. Recent thinking suggests that some combination of navigation functions [25] and juggling theory [15] may offer some help in solving such problems. This work is still very tentative.

Acknowledgement

It is a pleasure to acknowledge Dr. Elon Rimon who contributed many of the results reported here and was also responsible for generating all of the figures.

References

[1] J. R. Andrews and N. Hogan. Impedance control as a framework for implementing obstacle avoidance in a manipulator. In David E. Hardt and Wayne J. Book, editors, *Control of Manufacturing and Robotic Systems*, pages 243–251. A.S.M.E., Boston, MA, 1983.

[2] R. W. Brockett. On the computer control of movement. In *Proc. IEEE Conference on Robotics and Automation*, pages 534–540. IEEE, Philadelphia, PA., 1988.

[3] M. Bühler and D. E. Koditschek. From stable to chaotic juggling. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1976–1981, Cincinnati, OH, May 1990.

[4] M. Bühler and D. E. Koditschek. Stability analysis of a juggling robot. *IEEE Trans. Automatic Control*, (In preparation), 1991.

[5] M. Bühler, D. E. Koditschek, and P. J. Kindlmann. Planning and control of a juggling robot. *Int. J. Rob. Research*, (submitted), 1991 .

[6] M. Bühler, D. E. Koditschek, and P.J. Kindlmann. A family of robot control strategies for intermittent dynamical environments. *IEEE Control Systems Magazine*, 10:16–22, Feb 1990.

[7] Morris W. Hirsch and Stephen Smale. *Differential Equations, Dynamical Systems, and Linear Algebra*. Academic Press, Inc., Orlando, Fla., 1974.

[8] O. Khatib and J.-F. Le Maitre. Dynamic control of manipulators operating in a complex environment. In *Proceedings Third International CISM-IFToMM Symposium*, pages 267–282, Udine, Italy, Sep 1978.

[9] D. E. Koditschek and M. Bühler. Analysis of a simplified hopping robot. *Int. J. Rob. Research*, (to appear), 1991 .

[10] Daniel E. Koditschek. Exact robot navigation by means of potential functions: Some topological considerations. In *IEEE International Conference on Robotics and Automation*, pages 1–6, Raleigh, NC, Mar 1987.

[11] Daniel E. Koditschek. Applications of natural motion control. (to appear) *ASME Journal of Dynamics Systems and Measurement*, 1989.

[12] Daniel E. Koditschek. The Application of Total Energy as a Lyapunov Function for Mechanical Control Systems. In J. Marsden, Krishnaprasad, and J. Simo, editors, *Control Theory and Multibody Systems*, volume 97, pages 131–158. AMS Series in Contemporary Mathematics, 1989.

[13] Daniel E. Koditschek. Task encoding for autonomous machines: The assembly problem. In *Sixth Yale Workshop on Adaptive and Learning Systems*, pages 231–236, New Haven, CT, Aug 1990. Center for Systems Science, Yale University.

[14] Daniel E. Koditschek. Assembly: Another source of nonholonomy in robotics. In *Proc. American Control Conference*, page (to appear), Boston, MA, Sept. 1991. American Society of Control Engineers.

[15] Daniel E. Koditschek. Stabilizing feedback controllers for robot assembly tasks. In *Proc. IEEE Conf. on Decision and Control*, page (in preparation for submission), Brighton, UK, Dec. 1991. IEEE.

[16] Daniel E. Koditschek and Elon Rimon. Robot navigation functions on manifolds with boundary. *Advances in Applied Mathematics*, 11:412–442, 1990.

[17] Fumio Miyazaki and S. Arimoto. Sensory feedback based on the artificial potential for robots. In *Proceedings 9th IFAC*, Budapest, Hungary, 1984.

[18] V. V. Pavlov and A. N. Voronin. The method of potential functions for coding constraints of the external space in an intelligent mobile robot. *Soviet Automatic Control*, (6), 1984.

[19] E. Rimon and D. E. Koditschek. Exact robot navigation using cost functions: The case of spherical boundaries in $\mathbb{E}^n$. In *IEEE International Conference on Robotics and Automation*, pages 1791–1796, Philadelphia, PA, Apr 1988.

[20] E. Rimon and D. E. Koditschek. Exact robot navigation in topolgically simple but geometrically complicated environments. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1937–1943, Cincinnati, OH, May 1990.

[21] E. Rimon and D. E. Koditschek. The construction of analytic diffeomorphisms for exact robot navigation on star worlds. (to appear) *Transactions of the American Mathematical Society*, 1991.

[22] Elon Rimon and D. E. Koditschek. Exact robot navigation using artificial potential fields. *IEEE Trans. Robotics and Automation*, (In preparation), 1991.

[23] Elon F. Rimon. *Exact Robot Navigation Using Artificial Potential Functions*. PhD thesis, Yale University, New Haven, CT, Dec 1990.

[24] Jacob T. Schwartz and Micha Sharir. On the "Piano Movers" problem i. the case of a two-dimensional rigid polygonal body moving amidst polygonal barriers. Technical Report 39, N.Y.U. Courant Institute Department of Computer Science, New York, 1981.

[25] Louis W. Whitcomb and Daniel E. Koditschek. Navigation functions as a means of automatically generating plans and controllers for robot assembly tasks. In *Proc. IEEE International Workshop on Intelligent Robots and Systems*, page (in preparation for submission), Osaka, Japan, Dec. 1991. IEEE.