

Steering Nonholonomic Systems Using Sinusoids

Richard M. Murray* S. Shankar Sastry†

Electronics Research Laboratory
Department of Electrical Engineering and Computer Science
University of California Berkeley, CA 94720

Abstract

In this paper we investigate methods for steering systems with non-holonomic constraints between arbitrary configurations. Early work by Brockett derives the optimal controls for a set of canonical systems in which the tangent space to the configuration manifold is spanned by the input vector fields and their (first order) Lie brackets. Using Brockett's result as motivation, we derive suboptimal trajectories for systems which are not in canonical form and consider systems in which it takes more than one level of bracketing to achieve controllability. These trajectories use sinusoids at integrally related frequencies to achieve motion at a given bracketing level. Examples and simulation results are presented.

1 Introduction

Consider the problem of constructing a path $x(t) \in \mathbb{R}^n$ between a given x_0 and x_1 subject to k constraints which are linear in \dot{x} :

$$\omega_i(x)\dot{x} = 0 \quad i = 1, \dots, k$$

We assume the ω_i 's are smooth and linearly independent over the ring of smooth functions. Formally, these constraints are exterior differential one-forms on \mathbb{R}^n . This type of constraint arises in many mechanical systems, including systems with rolling contact and systems which conserve angular momentum. Specific examples of such systems are given in later sections.

Rather than use the machinery of exterior differential systems, we convert the problem to one in control theory. Let Δ be a distribution of dimension $m = n - k$ which is annihilated by the constraints. We represent this distribution with respect to a basis of vector fields:

$$\Delta = \text{span}\{g_1, g_2, \dots, g_m\} \quad g_i(x) \in \mathbb{R}^n$$

In coordinates, the constraint one-forms can be written as an $k \times n$ matrix and the g_i 's are a basis for the right null space of this matrix. The path planning problem can then be restated as finding an input function, $u(t) \in \mathbb{R}^m$ such that the control system

$$\dot{x} = g_1(x)u_1 + \dots + g_m(x)u_m \quad (1)$$

is driven from x_0 to x_1 . As a consequence of our assumptions on the ω_i 's, the g_i 's are also smooth and linearly independent.

The conditions for the existence of a path between two configurations is given by Chow's theorem. We let $[f, g]$ be the Lie bracket between two vector fields,

$$[f, g] = \frac{\partial g}{\partial x}f - \frac{\partial f}{\partial x}g$$

and define the involutive closure of a distribution Δ as the closure of Δ under Lie bracketing. Briefly, Chow's theorem states that if the involutive closure of the distribution associated with equation (1) spans \mathbb{R}^n at each configuration, the system can be steered between any two configurations. It is not apparent how the path can be explicitly constructed; in this paper we propose techniques for generating such paths.

*Supported in part by an IBM Manufacturing Fellowship

†Supported in part by NSF-PYI grant DMC84-51129

Some comments are in order with respect to the use of local control laws. It is tempting to adopt a local view of the system and achieve motion by stabilizing the system about the desired final configuration. To achieve global motion we could then stabilize about a sufficiently slow trajectory between configurations. Unfortunately, for nonholonomic systems, it can be shown that there is *no* smooth control law which stabilizes a point [3, 5]. Thus we are forced to consider more complicated approaches.

Related work

Many methods have been proposed for solving nonholonomic motion planning problems. Laumond performed early work in the context of path planning for mobile robots with a single nonholonomic constraint in the presence of obstacles [12]. Laumond and Siméon extended those results to steering a mobile robot with a trailer (two constraints) [13]. More recent work has been used for a mobile robot with bounded input constraints in addition to the nonholonomic constraint [10]. The approach presented in this second paper may be applicable to more general systems. Another general algorithm, developed by Barraquand and Latombe, can reportedly handle any set of non-holonomic constraints [2]. Sample paths are presented for a front-wheel drive car (two constraints) and a car pulling a trailer (three constraints). The paths generated are locally time-optimal, but can be computed only in sufficiently low-dimensional spaces. Yet another approach has been to train a neural net to park a car pulling a trailer [19].

Analytic techniques for studying nonholonomic motion have recently received attention. Li and Canny studied the motion of a fingertip rolling on an object without slipping [14]. This problem has also been investigated using some of the methods presented here [18]. Later work by Li and others studied a hopping robot flipping in mid-air by using conservation of angular momentum to construct paths on a reduced space [15]. Similar techniques have also been used for studying the motion of coupled rigid bodies and space manipulators [21, 16]. Lafferriere and Sussmann have recently proposed a method for general control systems using tools from geometric control theory and Lie algebra [11].

Our work is closed in spirit to that of Brockett [4]. Brockett considered a set of canonical systems in which the tangent space to the configuration manifold is spanned by the input vector fields together with their first order Lie brackets. He showed that the optimal controls (with respect to input cost) are sinusoids at integrally related frequencies. This paper works towards extending those results to both non-canonical systems and systems in which it takes more than one level of bracketing to achieve controllability.

Classification

We now develop some concepts which allow us to classify nonholonomic systems. A more complete treatment can be found in the work of Vershik [7, 23]. Basic facts concerning Lie algebras are taken from Varadarajan [22]. Let $\Delta = \text{span}\{g_1, \dots, g_m\}$ be the distribution associated with the control system (1). Define $G_1 = \Delta$ and

$$G_i = G_{i-1} + [G_1, G_{i-1}]$$

where

$$[G_1, G_{i-1}] = \text{span}\{[g, h] : g \in G_1, h \in G_{i-1}\}$$

The set of all G_i 's defines the *filtration* associated with a distribution. Each G_i is spanned by the input vector fields plus the vector fields formed by taking up to $i - 1$ Lie brackets. The Jacobi identity implies $[G_i, G_j] \subset [G_1, G_{i+j-1}] \subset G_{i+j}$.

A filtration is *regular* in a neighborhood U of x_0 if

$$\text{rank } G_i(x) = \text{rank } G_i(x_0) \quad \forall x \in U$$

We say a system is regular if the corresponding filtration is regular. If a filtration is regular, then at each step of the construction G_i either gains dimension or the construction terminates. If $\text{rank } G_{i+1} = \text{rank } G_i$ then G_i is involutive and hence $G_{i+j} = G_i$ for all $j \geq 0$. Clearly $\text{rank } G_i \leq n$ and hence if a filtration is regular, then there exists an integer $p < n$ such that $G_i = G_{p+1}$ for all $i \geq p + 1$. We refer to p as the *degree of nonholonomy* of the distribution.

For a regular system, a path exists between two arbitrary points in an open set $U \subseteq \mathbb{R}^n$ if and only if $G_p(x) = \mathbb{R}^n$ for all $x \in U$ (Chow's theorem). A system satisfying the conditions of the theorem is said to be *maximally nonholonomic*. This version of Chow's theorem is considerably weaker than the original version which holds for non-regular systems; a proof of the general case can be found in [8]. If a regular system is not maximally nonholonomic, then by Frobenius' theorem we can restrict ourselves to a manifold on which the system is maximally nonholonomic (see [22]).

It is also useful to record the dimension of each G_i . For a regular system, we define the *growth vector* $r \in \mathbb{Z}^{p+1}$ as

$$r_i = \text{rank } G_i$$

We define the *relative growth vector* $\sigma \in \mathbb{Z}^{p+1}$ as $\sigma_i = r_i - r_{i-1}$ and $r_0 := 0$. The growth vector for a system is a convenient way to represent information about the associated control Lie algebra. For a distribution with finite rank, the growth vector is bounded from above at each step. To properly determine this bound, we must determine the rank of G_i taking into account skew-symmetry and the Jacobi identity. A careful calculation [20] gives

$$\bar{\sigma}_i = \frac{1}{i} \left((\bar{\sigma}_1)^i - \sum_{j|i, j < i} \bar{\sigma}_j \right) \quad i > 1 \quad (2)$$

where $\bar{\sigma}_i$ is the maximum relative growth at the i^{th} stage and $j|i$ means all j such that j divides i . If $\sigma_i = \bar{\sigma}_i$ for all i , we say Δ has *maximum growth*.

2 First degree systems

Control systems in which the first level of brackets together with the input vector fields span the tangent space at each configuration arise in many areas. In classical mechanics, systems with growth vector $r = (n-1, n)$ are called contact structures [1]. A version of the Darboux theorem asserts that for these systems the corresponding constraint can be written as

$$dx_3 = x_2 dx_1$$

(using the notation of exterior differential forms). In \mathbb{R}^3 and using control system form, this becomes

$$\begin{aligned} \dot{x}_1 &= u_1 \\ \dot{x}_2 &= u_2 \\ \dot{x}_3 &= x_2 u_1 \end{aligned} \quad (3)$$

Brockett considered a more general version of this system [4]; we review his results here. Consider a control system as in equation (1) that is maximally nonholonomic with growth vector $(m, n) = (m, \frac{m(m+1)}{2})$. We would like to find an input $u(t)$ on the interval 0 to 1 which steers the system between an arbitrary initial and final configuration and minimizes

$$\int_0^1 |u|^2 dt$$

This problem is related to finding the geodesics associated with a singular Riemannian metric (Carnot-Carathéodory metric).

To solve the problem, Brockett considers a class of systems which have a special canonical form. An equivalent form, which is more useful for our purposes, is

$$\begin{aligned} \dot{x}_i &= u_i & i = 1, \dots, m \\ \dot{x}_{ij} &= x_i u_j & i < j \end{aligned} \quad (4)$$

We see that if $m = 2$, this is exactly the contact system (3). It can be shown that the input vector fields and their pairwise brackets span \mathbb{R}^n and hence the system is controllable with degree of nonholonomy equal to 1.

To find the optimal input between two points, we construct the Lagrangian

$$L(x, \dot{x}) = \sum_{i=1}^m \dot{x}_i^2 + \sum_{i,j} \lambda_{ij} (\dot{x}_{ij} - x_i u_j) \quad (5)$$

Here we have used the fact that $u_i = \dot{x}_i$. The λ_{ij} 's are the Lagrangian multipliers associated with the constraint imposed by the control system. Substituting equation (5) into the Euler-Lagrange equation

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{x}} - \frac{\partial L}{\partial x} = 0$$

it can be shown that the input must satisfy

$$u = e^{\Lambda t} u_0$$

where Λ is a constant skew-symmetric matrix. Thus the inputs are sinusoids at various frequencies. Unfortunately, even for very simple problems, determining Λ and u_0 given an initial and final configuration is very difficult.

A great deal of simplification occurs if we consider moving between configurations where $x_i(1) = x_i(0)$. In this instance the eigenvalues of Λ must be multiples of 2π and Brockett showed that the optimal inputs are sinusoids at integrally related frequencies, namely $2\pi, 2 \cdot 2\pi, \dots, \frac{m}{2} \cdot 2\pi$. This simplifies the problem tremendously and for many examples reduces the search to that of finding u_0 . We use this result to propose the following algorithm for steering systems of this type:

Algorithm

1. Steer the x_i 's to their desired values using any input and ignoring the evolution of the x_{ij} 's.
2. Using sinusoids at integrally related frequencies, find u_0 such that the input steers the x_{ij} 's to their desired values. By the choice of input, the x_i 's are unchanged.

The resulting trajectories are suboptimal but easily computable and have several nice properties which we will explore.

Example 1

We consider as an example a kinematic hopping robot, as shown in Figure 1. This example has been studied by Li, Montgomery and Raibert [15] using holonomy methods. We wish to reorient the body of robot while in midair and bring the leg rotation and extension to a desired final value. The kinematic equations of the robot (in center of mass coordinates) can be written as

$$\begin{aligned} \dot{\psi} &= u_1 \\ \dot{l} &= u_2 \\ \dot{\theta} &= \frac{m_l(l+1)^2}{1+m_l(l+1)^2} u_1 \end{aligned}$$

where we have used units such that the mass of the body is 1 and the length of the leg at zero extension is also 1. The last equation is a consequence of conservation of angular momentum. Expanding the equation using a Taylor series about $l = 0$

$$\dot{\theta} = \frac{m_l}{1+m_l} \dot{\psi} + \frac{2m_l}{(1+m_l)^2} l u_1 + o(l) u_1$$




Figure 1: A simple hopping robot. The robot consists of a leg which can both rotate and extend. The configuration of the mechanism is given by the angle of the body and the angle and length (extension) of the leg.




Figure 2: Nonholonomic motion for a hopping robot. Using sinusoidal inputs, the leg angle and extension return to their starting values but the body angle goes a net rotation.

This suggests a change of coordinates, $\alpha = \theta - \frac{m}{1+m} \psi$ to put the equations in the form

$$\begin{aligned} \dot{\psi} &= u_1 \\ \dot{l} &= u_2 \\ \dot{\alpha} &= \frac{2m}{(1+m)^2} l u_1 + o(l) u_1 = f(l) u_1 \end{aligned}$$

This equation has the same form locally as the canonical system in equation (4).

Using this as justification, we attempt to use our proposed algorithm to steer the full nonlinear system. Since we control the ψ and l states directly, we first steer them to their desired values. Then using sinusoids in the ψ and l inputs,

$$\begin{aligned} u_1 &= a_1 \sin \omega t \\ u_2 &= a_2 \cos \omega t \end{aligned}$$

we steer θ to its desired value. By construction, this last motion does not affect the final values of ψ and l . To include the effect of nonlinearity in the first vector field, harmonic analysis can be used. Since l is periodic, we expand f using its Fourier series,

$$f\left(\frac{a_2}{\omega} \sin \omega t\right) = \beta_1 \sin \omega t + \beta_2 \sin 2\omega t + \dots$$

Integrating $\dot{\alpha}$ over one period, only the first term in the expansion contributes to the net motion

$$\begin{aligned} \alpha\left(\frac{2\pi}{\omega}\right) &= \alpha(0) + \int_0^{\frac{2\pi}{\omega}} \left(\frac{a_1 \beta_1}{\omega} \sin^2 \omega t + \frac{a_1 \beta_2}{\omega} \sin \omega t \sin 2\omega t + \dots \right) dt \\ &= \alpha(0) + \frac{\pi a_1 \beta_1}{\omega^2} \end{aligned}$$

Figure 2 shows the trajectory for the last motion segment; ψ and l return to their initial values but α (and hence θ) experiences a net change. To compute the required input amplitudes, we plot β_1 as a function of a_2 and choose a_2 such that $\frac{\pi a_1 \beta_1}{\omega^2} = \theta_1 - \theta_0$. Using this procedure, we can (locally) steer between any two configurations.




Figure 3: Front wheel drive cart. The configuration of the cart is determined by the Cartesian location of the back wheels, the angle the car makes with the horizontal and the steering wheel angle relative to the car body. The two inputs are the velocity of the front wheels (in the direction the wheels are pointing) and the steering velocity. The rear wheels of the cart are always aligned with the cart body and are constrained to move along the line in which they point or rotate about their center.

3 Second and higher degree systems

We next consider systems in which the first level of bracketing is not enough to span \mathbb{R}^n . We begin by trying to extend the previous canonical form to the next higher level of bracketing. Consider a system which can be expressed as

$$\begin{aligned} \dot{x}_i &= u_i & i &= 1, \dots, m \\ \dot{x}_{ij} &= x_i u_j & i &< j \\ \dot{x}_{ijk} &= x_{ij} u_k & (\text{mod Jacobi identity}) \end{aligned} \quad (6)$$

Because Jacobi's identity imposes relations between certain brackets, not all x_{ijk} combinations are possible. This is analogous to limiting the x_{ij} 's to those for which $i < j$, reflecting skew-symmetry of the Lie bracket. Using the calculation in equation (2) shows that this system has relative growth vector $(m, \frac{m(m-1)}{2}, \frac{(m+1)m(m-1)}{3})$. Constructing the Lagrangian (with the same integral cost function) and substituting into the Euler-Lagrange equations does not result in a constant set of Lagrange multipliers. As a consequence, we cannot solve the optimal control problem in closed form.

We can however extend and apply our previous algorithm as follows:

Algorithm

1. Steer the x_i 's to their desired values. This causes drift in all other states.
2. Steer the x_{ij} 's to their desired values using integrally related sinusoidal inputs. If the i^{th} input has frequency w_i then x_{ij} will have frequency components at $w_i \pm w_j$. By choosing inputs such that we get frequency components at zero, we can generate motion in the desired states.
3. Use sinusoidal inputs a second time to move all previously steered states in a closed loop and generate motion only in the x_{ijk} directions. This requires careful choice of the input frequencies so that $w_i \pm w_j \neq 0$ but $w_i \pm w_j \pm w_k$ has zero frequency components.

Example 2

To illustrate the algorithm, we consider the motion of a front wheel drive car as shown in Figure 3. The kinematics of this mechanism can be written as

$$\begin{aligned} \dot{x} &= \cos \theta \cos \phi u_1 \\ \dot{y} &= \sin \theta \cos \phi u_1 \\ \dot{\phi} &= u_2 \\ \dot{\theta} &= \frac{1}{l} \sin \phi u_1 \end{aligned} \quad (7)$$

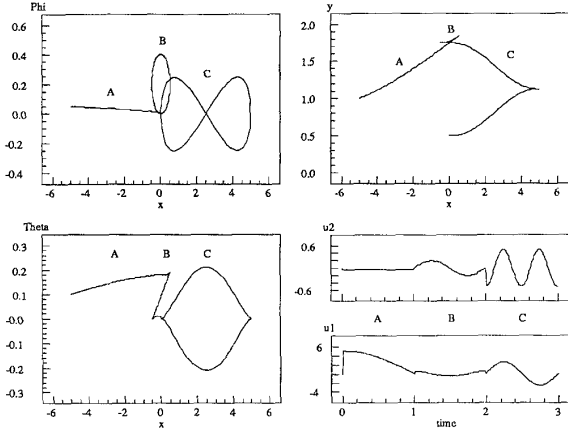


Figure 4: Sample Trajectories for a car. The trajectory shown is a three stage path which moves the unicycle from $x = -5, y = 1, \theta = 0.05, \phi = 1$ to $(0, 0.5, 0, 0)$. The first three figures show the states versus x ; the bottom right figures show the inputs as functions of time.

In this form, u_1 does not control any state directly. We use a change of coordinates and a change of input to put the equations in the form

$$\begin{aligned} \dot{x} &= v_1 & v_1 &= \cos \theta \cos \phi u_1 \\ \dot{\phi} &= v_2 & v_2 &= u_2 \\ \dot{\alpha} &= \tan \phi v_1 & \alpha &= \sin \theta \\ \dot{y} &= \frac{\alpha}{\sqrt{1-\alpha^2}} v_1 \end{aligned}$$

As before, the linear portion of the nonlinearities matches the canonical system and we can include the effects of the nonlinearities using Fourier series techniques.

An example of the algorithm applied to the car is given in Figure 4. The first portion of the path, labeled A, drives the x and ϕ states to their desired values using a constant input. The second portion, labeled B, uses a periodic input to drive θ while bringing the other two states back to their desired values. The last step brings y to its desired value and returns the other three states to their correct values. The Lissajous figures that are obtained from the phase portraits of the different variables are quite instructive. Consider the portion of the curve labeled C. The upper left plot contains the Lissajous figure for x, ϕ (two loops); the lower left plot is the corresponding figure for x, θ (one loop) and the open curve in x, y shows the increment in the y variable. The very powerful implication here is that the Lie bracket directions correspond to rectification of harmonic periodic motions of the driving vector fields and the harmonic relations are determined by the degree of the Lie bracket corresponding to the desired direction of motion. This point has also been made rather elegantly by Brockett [6] in the context of the rectification of mechanical motion.

The N -trailer problem

Consider the kinematic equations for a front wheel drive cart with N trailers as shown in Figure 5. The kinematic equations are

$$\begin{aligned} \dot{x} &= \cos \theta_0 \cos \phi u_1 \\ \dot{y} &= \sin \theta_0 \cos \phi u_1 \\ \dot{\phi} &= u_2 \\ \dot{\theta}_0 &= \frac{1}{l} \sin \phi u_1 \\ \dot{\theta}_i &= \frac{1}{d_i} \left(\prod_{j=1}^{i-1} \cos(\theta_{j-1} - \theta_j) \right) \sin(\theta_{i-1} - \theta_i) \cos \phi u_1 \end{aligned} \quad (8)$$

It may be verified that the system is completely controllable about the origin by using small angle approximations. It is also not difficult to see that with N trailers the degree of nonholonomy is $N + 2$ with relative growth vector $(2, 1, 1, \dots, 1)$. After redefining the inputs as we did for




Figure 5: Front wheel drive cart with trailers. The trailer configuration is described the the angle the trailer makes with the horizontal ϕ . The rear wheels of the trailer are fixed and constrained to move along the line in which they point or rotate about their center. The inputs to the system are the inputs to the tow car: the driving velocity (of the front wheels) and the steering velocity. This system is an example of a fourth degree system; higher degree systems can be generated by adding extra trailers.

the car, the system does not have the necessary triangular structure so that the harmonic analysis techniques of the previous section cannot be applied. In particular, it is not clear how to produce periodic motion in θ_1 since $\dot{\theta}_1$ depends on both θ_1 and θ_0 . However, it might be possible to use a change of coordinates and inputs to put the system into triangular form.

4 Extensions

A fundamental concept which we use repeatedly is that of a triangular system. A system is triangular if we can find a set of coordinates $h = (h^1, h^2, \dots, h^p) \in \mathbb{R}^{m_1 \times m_2 \times \dots \times m_p} = \mathbb{R}^n$ such that

$$\begin{aligned} \dot{h}^1 &= v & v &\in \mathbb{R}^{m_1} \\ \dot{h}^2 &= f^2(h^1)v \\ \dot{h}^3 &= f^3(h^1, h^2)v \\ &\vdots \\ \dot{h}^p &= f^p(h^1, \dots, h^{p-1})v \end{aligned}$$

The triangular form was necessary in our examples to insure that sinusoidal motion generated net movement in the i^{th} set of coordinates while leaving the previous sets of coordinates unchanged. Even if a system is not triangular, it may be possible to convert it to triangular form using a feedback transformation

$$\begin{aligned} h &= \phi(x) & \phi: \mathbb{R}^n &\rightarrow \mathbb{R}^n \\ v &= \beta(x)u & \beta(x) &\in \mathbb{R}^{m \times m} \end{aligned}$$

We saw an example of this in Section 3 where the car kinematics were converted to triangular form. Lafferriere and Sussmann have shown that the 1-trailer problem is also triangularizable [11].

Related to triangularity is nilpotentization. A distribution Δ is said to be nilpotent if there exists a basis for the distribution and an integer k such that all Lie products of degree $\geq k$ are identically zero. k is called the degree of nilpotency. Recent work with Montgomery has shown that a regular nilpotent system with maximum growth and degree of nilpotency equal to degree of nonholonomy can be triangularized [17]. Unfortunately, this result does not apply to any of the non-trivial examples we have presented in this paper since they are not nilpotent. However, it may be possible to convert a non-nilpotent system into a nilpotent one using a change of input $v = \beta(x)u$ [9]. We are currently studying the use of these techniques in order to apply sinusoidally generated motions to more general systems.

Another fundamental problem of great practical importance is obstacle avoidance. In many situations, such as parallel parking a car, obstacles play an important role in choosing a good trajectory. The use of sinusoidal inputs allows some freedom in shaping a trajectory based on the presence of obstacles. By varying the amplitude, phase, and number of cycles of the input sinusoids, different trajectories can be generated which result in the same net motion. Figure 6 shows three possible trajectories for parallel parking a car.




Figure 6: Alternative trajectories for parallel parking

Finally, we note that all of the trajectories we have studied have been fundamentally open loop. Errors in the initial conditions, model mismatch, and sensor noise will all degrade performance. A fundamental property of feedback control is robustness with respect to these disturbances. We are currently investigating methods for generating controllers for nonholonomic systems which respect the fundamental limitations mentioned in the introduction. An effective strategy may be to design controllers which perform some type of trajectory tracking rather than stabilization to a point.

We are also studying the effects of closed loop control on the planning stage. For example, it might be possible to simplify the non-holonomic path planning algorithm by considering simple classes of piecewise admissible paths. Feedback could then be used to decrease the initial configuration error present at the beginning of each segment. Another alternative is to plan the path of only a subset of the configuration variables. An example is a truck driver who concentrates only on the position of the end of the trailer when backing into a loading dock. The other state variables are determined by the control necessary to move the trailer in the desired direction. From the planning point of view, this method reduces the dimension of the space in which motions are generated and may result in computational savings.

Nonholonomic path planning represents a fusion of some of the newest ideas in control theory, classical mechanics, and differential geometry with some of the most challenging practical problems in robot motion planning. Furthermore, the class of systems to which the theory is relevant is broad: mobile robots, space-based robots, multifingered hands, and even such systems as a one-legged hopping robot. The techniques presented here indicate one possible method for generating efficient and computable trajectories for some of these nonholonomic systems in the absence of obstacles.

References

- [1] V.I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer-Verlag, second edition, 1989.
- [2] J. Barraquand and J-C. Latombe. On nonholonomic mobile robots and optimal maneuvering. In *4th International Symposium on Intelligent Control*, Albany, NY, 1989.
- [3] A. M Bloch and N. H. McClamroch. Control of mechanical systems with classical nonholonomic constraints. In *IEEE Control and Decision Conference*, pages 201–205, 1989.
- [4] R. W. Brockett. Control theory and singular Riemannian geometry. In *New Directions in Applied Mathematics*, pages 11–27. Springer-Verlag, New York, 1981.
- [5] R. W. Brockett. Asymptotic stability and feedback stabilization. In R.W. Brockett, R.S. Millman, and H.J. Sussman, editors, *Differential Geometric Control Theory*, pages 181–191. Birkhauser, 1983.
- [6] R. W. Brockett. On the rectification of vibratory motion. *Sensors and Actuators*, 20(1–2):91–96, 1989.
- [7] V. Gershkovich and A. Vershik. Nonholonomic manifolds and nilpotent analysis. *Journal of Geometry and Physics*, 5(3):407–452, 1988.
- [8] Robert Hermann and Arthur J. Krener. Nonlinear controllability and observability. *IEEE Transactions on Circuits and Systems*, AC-22:728–740, 1977.
- [9] Henry Hermes, Albert Lundell, and Dennis Sullivan. Nilpotent bases for distributions and control systems. *Journal of Differential Equations*, 55:385–400, 1984.
- [10] Paul Jacobs, Jean-Paul Laumond, and Michel Taix. A complete iterative motion planner for a car-like robot. In *Journées Geométrie Algorithmique*, INRIA, 1990.
- [11] G. Lafferriere and H.J. Sussmann. Motion planning for controllable systems without drift: a preliminary report. Technical Report SYCON-90-04, Rutgers Center for Systems and Control, June 1990.
- [12] J-P. Laumond. Finding collision-free smooth trajectories for a non-holonomic mobile robot. In *International Joint Conference on Artificial Intelligence*, pages 1120–1123, 1987.
- [13] J-P. Laumond and T. Siméon. Motion planning for a two degrees of freedom mobile robot with towing. In *IEEE International Conference on Control and Applications*, 1989.
- [14] Z. Li and J. Canny. Motion of two rigid bodies with rolling constraint. *IEEE Transactions on Robotics and Automation*, 6(1):62–71, 1990.
- [15] Z. Li, R. Montgomery, and M. Raibert. Dynamics and optimal control of a legged robot in flight phase. In *IEEE International Conference on Robotics and Automation*, pages 1816–1821, 1989.
- [16] R. Montgomery. Isoholonomic problems and some applications. *Communications in Mathematical Physics*, 128:565–592, 1990.
- [17] R. Montgomery, R. Murray, and S. Sastry. Triangulating differential equations on homogeneous nilpotent Lie groups. (work in progress).
- [18] R. Murray and S. Sastry. Grasping and manipulation using multifingered robot hands. In *Mathematical Question in Robotics, Lecture Notes*, Providence, RI, 1990. American Mathematical Society.
- [19] D. Nguyen and B. Widrow. Neural networks for self-learning control systems. *IEEE Control Systems Magazine*, 10(3):18–23, 1990.
- [20] J-P. Serre. *Lie Algebras and Lie groups*. W.A. Benjamin, New York, 1965.
- [21] N. Sreenath, Y.G. Oh, P.S. Krishnaprasad, and J.E. Marsden. The dynamics of coupled planar rigid bodies. part i: Reduction, equilibria and stability. *Dynamics and Stability of Systems*, 3(1 & 2), 1988.
- [22] V.S. Varadarajan. *Lie Groups, Lie Algebras, and Their Representations*. Springer-Verlag, 1984.
- [23] A. M. Vershik and V. Ya. Gershkovich. Nonholonomic problems and the theory of distributions. *Acta Applicandae Mathematicae*, 12:181–209, 1988.