

Progress in Desktop Robotics

Devin J. Balkcom and Matthew T. Mason
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, 15213, USA

Abstract

This paper describes our recent progress in Desktop Robotics. It begins with a description of the application domain, and then describes progress in a few areas: a mobile manipulator called the *mobipulator*, fast prototyping work culminating in the Palm Pilot Robot Kit, and work on time-optimal paths for diff drive robots.

1 Introduction

This paper surveys some recent work on “desktop robotics.” [1] The goal is a robotic system that handles common tasks on a typical office desktop, focusing primarily on the ability to manipulate documents and other common desktop objects.

Why desktop robotics? Robotics research is often shaped by an intended application domain. Obviously this is so for applied work, but it is true of theoretical work too. The research agenda in robotic manipulation is to a large extent determined by manufacturing applications. By developing an alternative application domain, we hope to develop new perspectives and insights on robotic manipulation. In particular, desktop robotics encourages: exploration of new manipulation hardware designs; use of mobile robots as manipulators; and exploration of human-robot interface issues.

2 Previous work

Since desktop robotics is a new application area there is not a lot of previous work to report, although additional details on desktop manipulation, the mobipulator, and time optimal paths for diff drives may be found at [1, 2, 3, 4, 5].

There has been a great deal of work exploring the connection between manipulation and locomotion, including Shakey [6], the JPL Cart [7], Donald et al’s [8] cooperating mobile manipulators, and

Romeo and Juliet [9] to name just a few examples. The JPL Cart, and Romeo and Juliet had arms, but the others just used mobile robots to push things, which is closer to the present approach. Perhaps closest of all is the Platonic Beast [10]. Just as the mobipulator uses its wheels for either locomotion or manipulation, the Platonic Beast used its limbs for either locomotion or manipulation.

There has also been relevant work in manipulation. The present approach is an example of “non-prehensile manipulation” or “graspless” manipulation. Perhaps the closest previous work is the business card manipulation work of Kao and Cutkosky [11], which addressed manipulation of laminar objects by fingers pressing down from above. See Lynch [12] for a more detailed survey of previous work on nonprehensile manipulation.

Previous work on optimal control of diff drive robots has assumed bounds on acceleration rather than on velocity; for example see papers by Reister and Pin [13] and Renaud and Fourquet [14]. For steered vehicles the original work was by Dubins [15] and by Reeds and Shepp [16]. Our work is an extension of optimal control techniques developed for steered vehicles by Souères, Boissonnat, Laumond, Sussman, Tang, and others [17, 18, 19].

3 The mobipulator

This section describes our work with the “mobipulator,” which is a contraction of “mobile manipulator.” The goal of the mobipulator is to explore the idea of using wheels for manipulation.

3.1 Mobipulator design

The mobipulator is about 10 cm by 10 cm, and has four independently powered wheels. A thin strip of steel connects the front and rear axles; the strip is flexible enough to act as a suspension, allowing the weight to be evenly distributed among the four wheels. The motors are DC servomotors

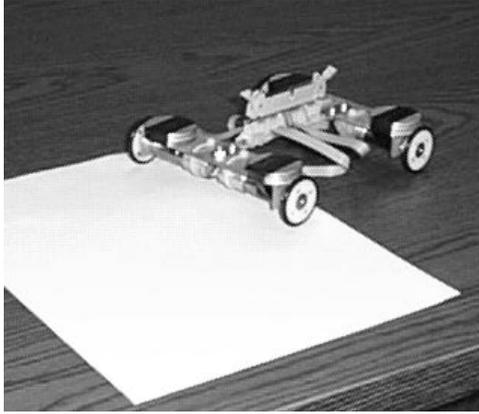


Figure 1: The mobipulator

with optical encoders. The machine is tethered—power electronics, IO, and computation are all off-board. There is no steering mechanism, although skid steering is possible.

The mobipulator can navigate an uncluttered desktop. It uses its wheels both for locomotion and for manipulation. To move a piece of paper it typically rolls halfway onto the paper. It then uses its front wheels to manipulate the paper while it uses its rear wheels to locomote. To move a cylinder it must first be placed with its front wheels on the cylinder. Then it runs its front wheels backward as it runs its rear wheels forward, much as a human might roll a barrel.

A mobile robot has a workspace much larger than itself. This gives a small mobile manipulator on the desk several advantages: the robot is small, inexpensive, unobtrusive, and intrinsically safe. The current design also has some compensating disadvantages: it is unable to navigate rough terrain, accurate position estimation requires additional sensing, and manipulation of even simple objects is challenging.

3.2 Theoretical results

We have focussed our work so far on the task of moving a piece of paper around on the desktop. We model both the mobipulator and paper as rigid bodies in the plane.

3.2.1 Contact modes and configuration space: The effects of the controls depend on the location of the mobipulator’s wheels relative to the piece of paper. The simplest approach is to assume the paper is rigid and to neglect the case of a wheel at the edge of the paper. Since there are four wheels and each can be either on or off of the paper, there

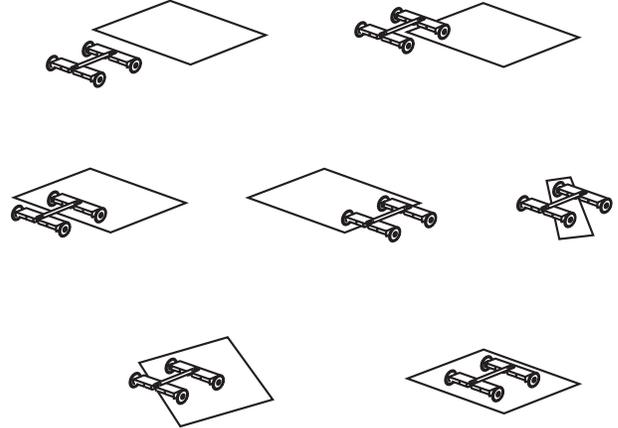


Figure 2: Seven of the sixteen contact modes

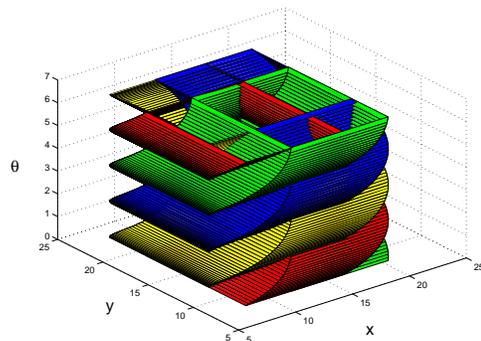


Figure 3: Boundaries of contact modes in configuration space. Coordinates are location and orientation of robot relative to paper.

may be up to sixteen contact modes, depending on the shape of the paper. Figure 2 shows the seven important cases; the other cases are symmetric.

The configuration of the robot relative to the paper determines which wheels are in contact with the paper. We can construct the configuration space of the robot relative to a piece of paper. We choose the reference point to be the center of the mobipulator. Then we choose a wheel, and construct the configuration space boundary of the paper. Since we assume the wheel makes a point contact, each slice of the boundary with θ held fixed will have the same shape as the paper; for configurations on this boundary, the selected wheel is touching the edge of the paper. The boundary encloses a region for which the wheel is on the paper; we repeat the process for each wheel and intersect the regions to find the regions of configuration space corresponding to various contact modes. Figure 3 shows an example for a rectangular piece of paper.

3.2.2 Dual diff drive mode: One of the most interesting modes is “dual diff drive” mode, for which the robot has either only its front or only its rear wheels touching the paper. We refer to the wheels on the paper as the “hands,” and to the wheels on the desktop as “feet.” The feet then control the system as a diff drive, and the hands control the paper as a diff drive. (A diff drive robot is like a wheel chair, with two independently driven coaxial wheels.) Choosing appropriate state variables, we derive the following kinematic model:

$$\begin{bmatrix} \dot{x}_f \\ \dot{y}_f \\ \dot{\theta}_f \\ \dot{x}_h \\ \dot{y}_h \\ \dot{\theta}_h \end{bmatrix} = \begin{bmatrix} \cos \theta_f & 0 & 0 & 0 \\ \sin \theta_f & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos \theta_h & 0 \\ 0 & 0 & \sin \theta_h & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_f \\ \omega_f \\ v_h \\ \omega_h \end{bmatrix} \quad (1)$$

where v_f and ω_f are the velocity and angular velocity of the robot relative to the desktop, and v_h and ω_h are the velocity and angular velocity of the robot relative to the paper.

It is quite easy to verify that this mode is nonholonomic by taking Lie brackets. The other modes also have either holonomic or nonholonomic constraints.

3.2.3 Lower dimensional goals: local paper moving: The nonholonomic constraints of the dual diff drive mode make the planning and control problem difficult. However, if we wish only to move the paper to a new configuration and there are no obstacles, the following approach is useful. Let the state of the paper be (x_p, y_p, θ_p) . Then we can use some geometry to derive a matrix $A(q)$, a function of the state, such that for some desired $(\dot{x}_p, \dot{y}_p, \dot{\theta}_p)$ we can derive required controls $(v_f, \omega_f, v_h, \omega_h)$. We have an equation of the form:

$$\begin{bmatrix} v_f - v_h \\ \omega_f \\ \omega_h \end{bmatrix} A(q) = \begin{bmatrix} \dot{x}_p \\ \dot{y}_p \\ \dot{\theta}_p \end{bmatrix} \quad (2)$$

We can then derive a simple linear control law using $(\dot{x}_p, \dot{y}_p, \dot{\theta}_p)$ as the controls to drive (x_p, y_p, θ_p) to some desired value, assuming $A(q)$ is defined over the region of interest.

3.3 Experience with a planner

We have implemented a planner based on one suggested by Barraquand and Latombe in [20] for a steered car. The planner simulates a few constant controls for a small duration to generate discrete

actions. A heuristically guided search is then used to find a path to a configuration in the neighborhood of the goal. One reason for choosing this style of planner is that it can generate fast plans. Although we are not really concerned with speed, the fastest plans are often interesting and unintuitive, involving the robot first manipulating the paper with its front wheels, then driving across the paper and manipulating with rear wheels.

We were not entirely satisfied with our implementation of this planner. We modelled the system with six degrees of freedom; the discretization necessary for reaching the goal with reasonable accuracy make the search in this space intractable. We therefore split the task into two three-DOF tasks: one task to reach a configuration of the robot relative to the paper, and a second task to reach a configuration of the robot relative to an origin on the desk. Even so, the planner was slow and required a lot of memory. There are a number of factors which contributed to this:

- The planner by Barraquand and Latombe attempts to minimize time and number of control switches. Unfortunately, this means that to be accurate we must keep track of an additional state variable, the number of switches. This means the search is in four dimensions rather than three.
- The chosen discretization of the controls and θ depends on the desired translational accuracy of the goal, and the distance of the goal from that start. For example, consider that an obvious plan for a diff drive is to turn towards the goal and then drive to it. If the goal is far away, much more precision is necessary to turn to exactly the required angle. This problem is less evident for the steered car that Barraquand and Latombe studied, which cannot turn in place.
- Manipulation goals may be much further away than locomotion goals of the type studied by Barraquand and Latombe.

3.4 Experimental results

Our analysis has considered a kinematic model, and we have assumed the paper acts as a rigid body. However, the manipulator will be much more capable if it can exploit knowledge of dynamics or the flexibility of the paper. In the lab we have tried a number of interesting manipulation techniques.

For example, if all four wheels are on the pa-

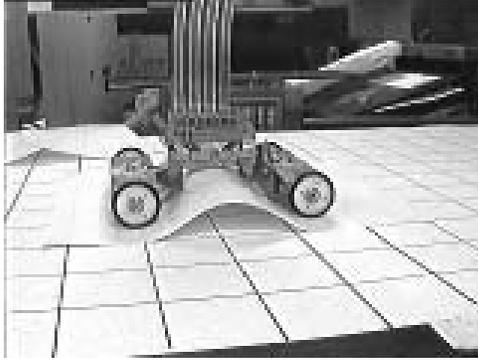


Figure 4: Inchworm mode

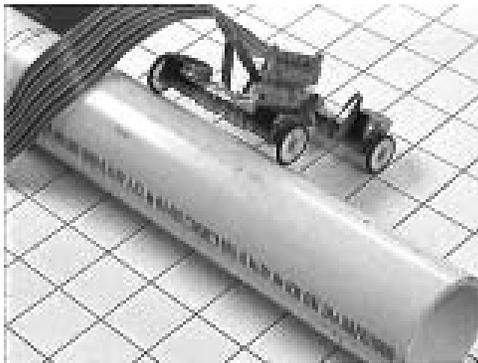


Figure 5: Mobipulator rolling PVC pipe

per, the mobipulator can spin its front wheels and rear wheels in opposite directions. This creates a “hump” in the paper. The robot can inch the paper forward under the front wheels by a very quick acceleration of the wheels. By alternating between front wheels and rear wheels, the paper can move forward while staying humped. (See figure 4.) Although “inchworm” actions of this type appear impractical due to their slow speed, they might be useful for small adjustments or when obstacles inhibit access.

Another interesting action is the “running slide.” If the mobipulator drives quickly onto the paper and then applies the brakes, the momentum of the system will cause it to slide forwards.

The mobipulator can also roll a cylinder by rolling the front wheels forward as the rear wheels roll backward. For large cylinders the motion is quite stable. Figure 5 shows the mobipulator rolling PVC pipe.

3.5 New perspectives?

The motivation for designing a new application domain was to develop new perspectives on manip-



Figure 6: The Palm Pilot Robot Kit

ulation. The study of manipulation using wheels is a new area for us, and the new perspective has had good results. One result is theoretical in nature: dual diff drive mode of the mobipulator motivated us to study the optimal trajectories for differential drive robots. Mobile robots typically have high mass relative to their motor torque, so previous studies have considered models of the diff drive with bounds on the acceleration of the wheels; unfortunately the problem has been too difficult to solve analytically. In our problem, the paper acts as a diff drive. The paper has low mass, so it is natural to assume bounds on wheel velocities rather than on accelerations. The model is simpler, and we were able to find the optimal trajectories analytically. This work is described later in the paper. First we describe some work on fast prototyping.

4 Fast prototyping

One of the goals for the new application domain was to encourage research in new manipulation hardware concepts. Inspired in part by Inaba’s work on remote-brained robotics,[21] we set out to develop the lab resources and skills to quickly prototype and evaluate new designs.

The materials and techniques are well known to the robotics research community and to the hobby robotics movement as well. We rely heavily on model airplane equipment, including the radio transmitters and receivers as well as the servomotors. For structural elements we use modular kits such as Lego and Meccano, plastics, and sheet metal. Control electronics included Palm Pilots, the Pontech serial IO card, HC11 microcontrollers, PIC based cards. For sensors we used simple IR

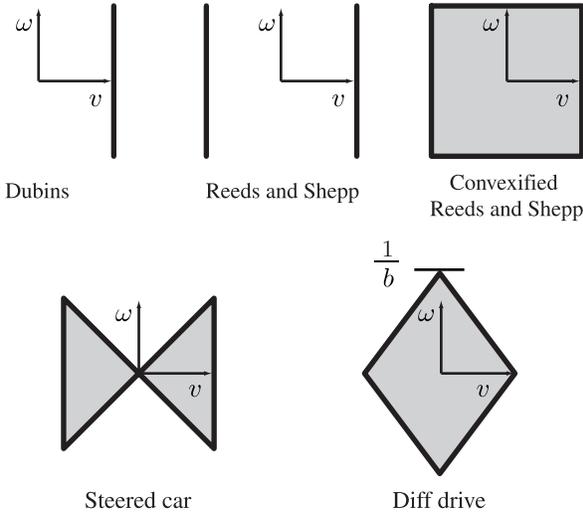


Figure 7: Bounded velocity models of mobile robots

proximity detectors, and the Sharp IR range sensor.

The main result of this effort was a robot that used three omnidirectional wheels to locomote in the plane. Through the efforts of several collaborators this device evolved into the Palm Pilot Robot Kit, which is documented on the web at <http://www.cs.cmu.edu/~pprk>. It has been spectacularly successful with the popular press, although it has yet to demonstrate any useful desktop manipulation. Its main value is educational—it is an easy way for a novice to learn about simple robotic systems.

Our fast prototyping infrastructure has been moderately successful, in that we can put working prototypes together in a matter of hours. The prototype Palm Pilot Robot Kit (actually using a model airplane radio receiver rather than a Palm Pilot) took less than one hour from concept to working model.

5 Time optimal paths

The motion planning algorithm described earlier requires a small discrete set of actions to generate candidate paths. One way to choose this set is to use Pontryagin’s Maximum Principle to identify the time optimal motions. Since the mobipulator is closely related to a diff drive, this led us to the problem of identifying the time optimal paths for a diff drive robot, assuming the wheel velocities are independently bounded, with no bounds on acceleration.

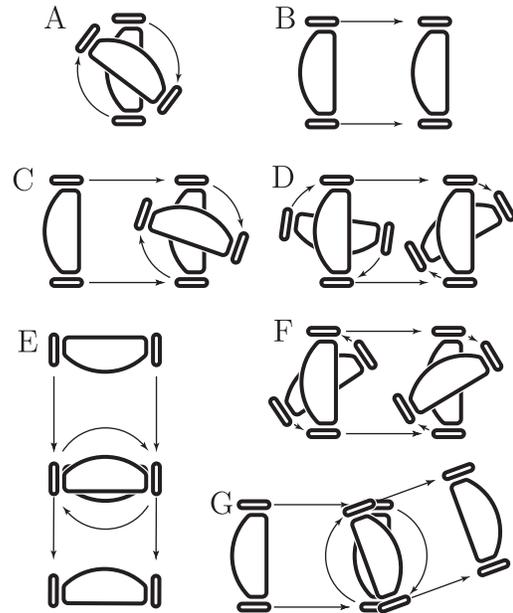


Figure 8: The seven simplest optimal trajectory classes.

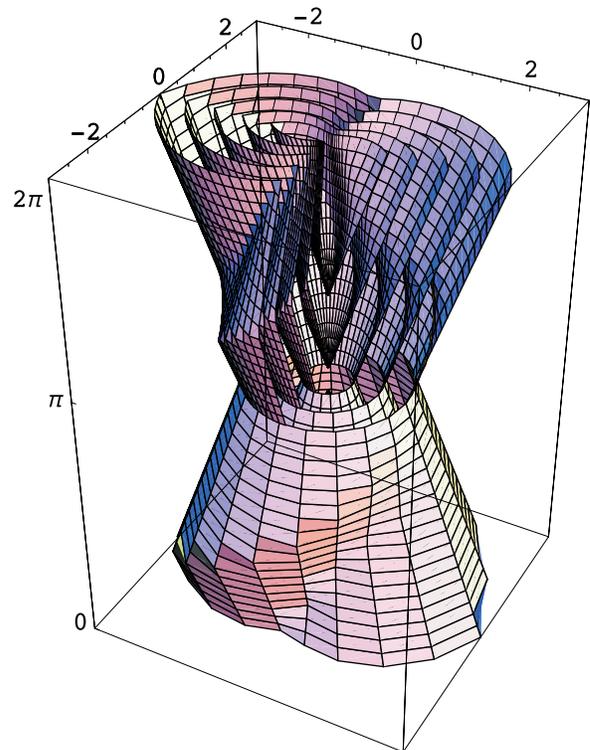


Figure 9: Reachable configurations in normalized time 1, 1.5, 2, 2.5, and 3.

The time optimal paths for a steered car are well known, and it might seem that a diff drive is equivalent to a steered car with zero turning radius. There is a difference though. For the usual model of a steered car, the time to turn is proportional to the turning radius, so with zero turning radius the car would turn in place instantaneously. Figure 7 compares the feasible controls of the diff drive with that of closely related models, including the steered car.

Pontryagin's Maximum Principle is easily applied to the problem, and the resulting equations can be solved analytically. Further analysis completely characterizes the time optimal paths. A time optimal path is composed of straight lines and turns in place. There are forty different combinations, arranged in nine different symmetry classes. Examples for the simplest seven classes are shown in figure 8. Figure 9 shows the configurations the robot can reach in normalized time 1, 1.5, 2, 2.5, and 3.

6 Future work

There are some glaring deficiencies in the present designs. The mobipulator cannot navigate a cluttered desk, nor can it manipulate more massive or complex objects on a desktop. We are continuing to prototype different designs for mobile manipulators to address these problems. We would also like to explore competitive approaches, such as a fixed base manipulator. Or perhaps we are asking the wrong question. Instead of "what robot do you want on your desk?" perhaps we should be asking "what do you want instead of a desk?"

We are also extending the fast prototyping work to include control, sensing, and programming of the robots. We are building a control and programming system to make it easy to try different controls, different sensory systems, and to program new behaviors. We hope to dramatically expand the experience base, and explore several new manipulation techniques with the mobipulator and other prototypes.

Acknowledgements

The Desktop Robotics domain and the initial work on the mobipulator was a collaboration with Dinesh Pai, Daniela Rus, Jon Howell, Lee Taylor, and Mike Erdmann. The Palm Pilot Robot Kit was a collaboration with Greg Reshko, Illah Nourbakhsh,

Garth Zeglin, Joshua Mouch, Angela Demke, Allison Bruce, and Chris Hardouin. We would like to thank Jean Paul Laumond for guidance on the optimal control work. This research was supported by the National Science Foundation under grant IIS-9900322, and by a DOE graduate fellowship.

References

- [1] Matthew T. Mason, Dinesh K. Pai, Daniela Rus, Jon Howell, Lee R. Taylor, and Michael A. Erdmann. Experiments with desktop mobile manipulators. In Peter Corke and James Trevelyan, editors, *Experimental Robotics VI: The Sixth International Symposium*, pages 37–46. Springer, 2000.
- [2] Matthew T. Mason, Dinesh K. Pai, Daniela Rus, Lee R. Taylor, and Michael A. Erdmann. A mobile manipulator. In *International Conference on Robotics and Automation*, pages 2322–2327. IEEE, 1999.
- [3] Devin J. Balkcom and Matthew T. Mason. Time optimal trajectories for bounded velocity differential drive robots. In *IEEE International Conference on Robotics and Automation*, 2000.
- [4] Devin J. Balkcom and Matthew T. Mason. Extremal trajectories for bounded velocity differential drive robots. In *IEEE International Conference on Robotics and Automation*, 2000.
- [5] Devin J. Balkcom and Matthew T. Mason. Graphical construction of time optimal trajectories for differential drive robots. In Bruce R. Donald, Kevin M. Lynch, and Daniela Rus, editors, *Algorithmic and Computational Robotics: New Directions*, The Fourth Workshop on the Algorithmic Foundations of Robotics, pages 377–390, Natick, Massachusetts, 2001. A. K. Peters.
- [6] Nils J. Nilsson. Shakey the robot. Technical Report 323, SRI International, 1984.
- [7] Alan M. Thompson. The navigation system of the jpl robot. In *Fifth International Joint Conference on Artificial Intelligence*, pages 749–757, 1977.
- [8] Bruce Randall Donald, James Jennings, and Daniela Rus. Analyzing teams of cooperating mobile robots. In *Proceedings 1994 IEEE International Conference on Robotics and Automation*, 1994.
- [9] Oussama Khatib, K. Yokoi, K. Chang, D. Ruspini, R. Holmberg, A. Casal, and A. Baader. Force strategies for cooperative tasks in multiple mobile manipulation systems. In Georges Giralt and Gerhard Hirzinger, editors, *Robotics Research:*

The Seventh International Symposium, pages 333–342, 1996.

[10] Dinesh K. Pai, Roderick A. Barman, and Scott K. Ralph. Platonic beasts: a new family of multilimbed robots. In *Proceedings 1994 IEEE International Conference on Robotics and Automation*, pages 1019–1025, 1994.

[11] Imin Kao and Mark R. Cutkosky. Dextrous manipulation with compliance and sliding. In Hirofumi Miura and Suguru Arimoto, editors, *Robotics Research: the Fifth International Symposium*. Cambridge, Mass: MIT Press, 1990.

[12] Kevin M. Lynch. *Nonprehensile Robotic Manipulation: Controllability and Planning*. PhD thesis, Carnegie Mellon University, March 1996. CMU-RI-TR-96-05.

[13] David B. Reister and Francois G. Pin. Time-optimal trajectories for mobile robots with two independently driven wheels. *International Journal of Robotics Research*, 13(1):38–54, February 1994.

[14] Marc Renaud and Jean-Yves Fourquet. Minimum time motion of a mobile robot with two independent acceleration-driven wheels. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 2608–2613, 1997.

[15] L. E. Dubins. On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79:497–516, 1957.

[16] J. A. Reeds and L. A. Shepp. Optimal paths for a car that goes both forwards and backwards. *Pacific Journal of Mathematics*, 145(2):367–393, 1990.

[17] P. Souères and J.-D. Boissonnat. Optimal trajectories for nonholonomic mobile robots. In J.-P. Laumond, editor, *Robot Motion Planning and Control*, pages 93–170. Springer, 1998.

[18] Philippe Souères and Jean-Paul Laumond. Shortest paths synthesis for a car-like robot. *IEEE Transactions on Automatic Control*, 41(5):672–688, May 1996.

[19] Héctor Sussmann and Guoqing Tang. Shortest paths for the reeds-shepp car: a worked out example of the use of geometric techniques in nonlinear optimal control. SYCON 91-10, Department of Mathematics, Rutgers University, New Brunswick, NJ 08903, 1991.

[20] Jérôme Barraquand and Jean-Claude Latombe. Nonholonomic multibody mobile robots: Controllability and motion planning in the

presence of obstacles. *Algorithmica*, 10:121–155, 1993.

[21] Masayuki Inaba. Remote-brained robotics: Interfacing ai with real-world behaviors. In Takeo Kanade and Richard Paul, editors, *Robotics Research: The Sixth International Symposium*, pages 335–344, Cambridge, USA, 1993. The International Foundation for Robotics Research.