

#779

AN ALGORITHMIC APPROACH TO SOME PROBLEMS IN TERRAIN NAVIGATION

(Short Title: An Algorithmic Approach to Terrain Navigation)

Presented at the Oxford Workshop on Geometric Reasoning
Oxford University, June 30 – July 3, 1986

Joseph S. B. Mitchell*

School of Operations Research and Industrial Engineering
Cornell University
Ithaca, NY 14853

Abstract

Recent advances in the field of computational geometry have provided efficient algorithms for a variety of shortest path problems. Many problems in the field of terrain navigation can be cast as optimal path problems in a precise geometric model. With such a model one can develop and analyze algorithms for the solution of the original problem and can gain insights into how to design more efficient heuristics to deal with more complex problems. We examine the path planning problem in which we are given a “map” of a region of terrain and we are expected to find optimal paths from one point to another. This, for example, is a task which must be done repeatedly for the guidance of an autonomous vehicle. We examine how to formulate some path planning problems precisely, and we report algorithms to solve certain special cases.

Key Words: computational geometry, shortest paths, terrain navigation, autonomous vehicles, route planning

1. Introduction

One of the basic tasks required of an autonomous vehicle is to move from one location to another. Usually, the robot has some knowledge (which we refer to as a *map*) of the terrain on which it is navigating, although it may initially have none. It should try to exploit as much of this knowledge as possible in order to make the best decisions about its path. Two types of navigation problems immediately arise: How does the robot *locate* its position in the map, and how does one use the map and the current location of the robot to *plan*

* Partially supported by a grant from Hughes Research Laboratories, Malibu, CA, and by National Science Foundation grant IRI-8710858.

a route to the goal? In this paper, we concentrate on the latter type of problem. We use the word “terrain” to refer to both indoor and outdoor environments for a vehicle.

The problem of location involves using sensory data to compute the coordinates of the robot. This can be done only to within some error bounds; however, we will assume that the errors involved are small enough in comparison with the scale of the planning problem that we can assume the exact position of the robot is known at each point in time and that the robot can be controlled exactly to follow a specified trajectory. The issue of location uncertainty is very important, and it is an area of continuing research.

In planning paths through the map, we typically wish to minimize some objective function which specifies the *cost* of motion along the path (the *path length*). Depending on how the terrain is modeled and what the optimization criteria are, this results in various versions of the *shortest path problem*. The shortest path problem for a mobile robot was one of the early problems addressed by researchers in artificial intelligence (see Nilsson [Ni] and Moravec [Mo]).

In this paper we discuss some of the issues that go into selecting a representation of a map, and we report some of the recent results from the field of *computational geometry* which are of use in solving terrain navigation problems. Computational geometry and algorithmic motion planning have recently erupted as important areas of research in computer science. We certainly cannot do justice here to the field of motion planning, so we refer the interested reader to the superb new books of Schwartz and Yap [SY] and Schwartz et al. [SSH], which collect together many of the landmark papers of this exciting field. We also refer the reader to the award-winning thesis of Canny [Ca]. Much of our discussion uses concepts from the field of computational geometry. We refer the interested reader to three recent books on the subject: Edelsbrunner [Ed], Mehlhorn [Meh], and Preparata and Shamos [PrS].

While many of the problems of terrain navigation have been addressed before from a heuristic or approximation point of view, we feel that it is important to formalize some of the models and to discuss the computational complexity of exact algorithms. We have found that an understanding of the underlying geometry can lead to new heuristics and faster algorithms for a variety of problems.

How do we go about representing varied terrain? What representations yield particularly natural or efficient algorithms for path planning? What do shortest paths “look like”; that is, can we understand enough about the local optimality criteria to cut down the search for optimal paths?

We begin addressing some of these issues in the following sections, where we describe the obstacle avoidance problem, the *Discrete Geodesic Problem*, the *Weighted Region Problem*, and some special cases of these problems that arise naturally in finding shortest paths through terrain. Throughout our discussions, we concentrate on the fundamental algorithmic issues of computing paths, so we are interested in worst-case asymptotic running times. Where appropriate, we will allude to the practicality of the algorithms and approaches we suggest.

2. Obstacle Avoidance in the Plane

The simplest surface to model is a plane, and the simplest nontrivial terrain map on a plane is a binary partitioning of the plane into regions that are obstacles and regions that are traversable by the vehicle. Usually, the robot can be modeled as a point, which is a reasonable assumption if it is small in comparison with the dimensions of the terrain map. Also, by *configuration space* techniques, many problems of moving a non-point robot among obstacles can be reduced to the case of moving a point (possibly in higher dimensions), by “growing” the obstacles in an appropriate way, as in Lozano-Pérez and Wesley [LW]. The problem of planning a shortest Euclidean path from one point to another is then the usual *obstacle-avoidance shortest path problem*. This problem has been studied by many researchers over the last twenty years. A survey of many algorithms is given in Mitchell and Papadimitriou [MP1] or Mitchell [Mi2].

Some recent work of Papadimitriou and Silverberg [PaS] and O’Rourke [O’R] has started to address the tough problem of *shortest* paths for a non-point non-circular body in the plane. While this can be mapped into a shortest path problem in three-dimensional configuration space, calculating shortest paths in three dimensions has been shown to be quite difficult, as we will mention later.

There are two standard approaches to modeling the terrain for the obstacle avoidance problem. The first subdivides the plane into small regular pieces (“pixels”) and labels each piece as obstacle or nonobstacle. If all of the pieces are squares of exactly the same size, then we get the standard grid tessellation, and we can represent the map as a binary array. A measure of the size of the problem instance is given by n , the total number of pixels. Of course, n depends on the resolution of the grid used to approximate the set of obstacles. Picking a very small pixel (very high resolution) allows a better model of the obstacle space, but it costs us significantly in problem size. If using a regular grid, we may be forced to have pixels smaller than a certain size in order to assure that skinny passageways between obstacles or tiny little obstacles appear in the map. Other subdivision representations exist which are hierarchical in nature, such as the *quadtree* (see Samet and Webber [Sa] or Kambhampati and Davis [KD]). These schemes have the advantage that they conglomerate groups of contiguous small squares into single larger squares.

Another standard approach to modeling terrain which is obstacle or non-obstacle is to give a description of a polygonal space in terms of its boundary representation. Obstacles are given as a list of k simple polygons, each represented by a doubly-linked list of vertices (each vertex just being a pair of coordinates, either integer or real). Usually, the obstacles are assumed to be disjoint. Let n be the total number of vertices of all polygonal obstacles. We refer to n as the “complexity of the scene”, as it is proportional to the number of bytes of storage needed to represent the map. To be more precise, though, the complexity of the scene should be the pair (n, k) , since some of the complexities of algorithms depend not only on n , but also on k . (Of course, k is always less than or equal to n , but it may be significantly less than n , in which case it is important to know precisely how the running times depend on k .) The objective of the algorithmic study of the shortest path problem is to determine an algorithm whose worst-case running time

is a low-degree polynomial in n and k . Most work in computational geometry focuses on *worst-case* running times; however, it is an important issue (not addressed here) to consider the *average-case* running times of shortest path algorithms.




Figure 1(a). Polygonal obstacles and a shortest path.




Figure 1(b). Digitized representation.

An example of a set of obstacles is given in Figure 1. In Figure 1(a), we show the representation of the obstacles as a set of $k = 2$ polygons with $n = 8$ vertices and the shortest path from start to goal. Then, in Figure 1(b) we show the digitized representation on a grid of 16-by-16 pixels ($n = 256$). In Figure 1(c) we show the same set of obstacles represented as a quadtree, where the free space is also subdivided into a quadtree representation. The total size of the representation in this case is $n = 124$.

There are other methods of modeling the terrain in the obstacle-avoidance problem. The obstacles may be given as a list of circles, ellipses, or other non-polygonal objects. The shortest path problem among circular or elliptical obstacles has been addressed by Baker [Bak], Chew [Ch], and Mitchell [Mi2]. Souvaine [So]

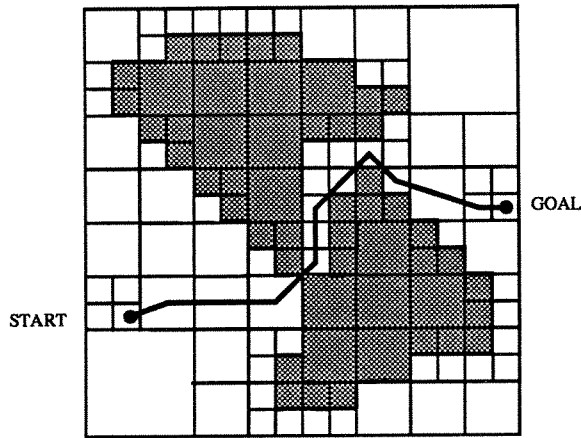


Figure 1(c). Quadtree representation.

introduced the notion of *splinegons* as a means of describing simple closed curved figures, and she describes many generalizations of standard computational geometry algorithms to the case of splinegons. A different approach to modeling obstacle maps is given by Brooks [Br] and Brooks and Lozano-Pérez [BL], who give a decomposition of free space into generalized cylinders. Their (heuristic) method is similar to the retraction method of motion planning (see Schwartz and Yap [SY]).

Shortest paths according to the Euclidean metric will be “taut-string” paths: imagine threading a string from START to GOAL among the obstacles, then pulling the string until it is taut. The question really comes down to how the string should be threaded. One cannot afford to try all of the exponentially many possible threadings.

The algorithm one uses to find shortest paths among obstacles depends on the representation used in the map. For the case of a binary grid representation, the search for shortest paths is easy and straightforward. Shortest paths are found by restricting the robot’s motion to connect centers of adjacent free squares. Then, we can consider the graph (so-called *grid graph*) whose nodes are the free space squares and the edges connect adjacent free squares. The degree of a node is thus at most 8, so the graph has at most n nodes and no more than $8n$ edges. Lengths are assigned to edges in the obvious way: diagonal edges are of length $\sqrt{2}$ times the size of a square, and horizontal and vertical edges are of length equal to the size of a square. We can then search the graph for shortest paths by applying the A^* algorithm (Nilsson [Ni]) or the Dijkstra [Di] algorithm (which is “uninformed” A^*). The worst-case running time of Dijkstra’s algorithm (and hence of A^*) is $O(|E| + |V| \log |V|)$ for a graph with $|E|$ edges and $|V|$ nodes (Fredman and Tarjan [FT]); thus, the worst-case running time of the shortest path algorithm is $O(n \log n)$ for a grid with n pixels. Note that for A^* there is an obvious admissible heuristic function, $h(v)$, for node v , namely, the straight-line Euclidean distance from v to the goal.

Figure 1(b) shows the path from START to GOAL which results from finding the shortest path through

the grid graph. Note that the path is not shortest according to the Euclidean metric. This is a common problem with the grid graph approach – it introduces a “digitization bias”, a product of requiring paths to stay on the grid graph. Within the same time bounds ($O(n \log n)$), an algorithm of Keirse and Mitchell [KM] can be run which attempts to correct for this digitization bias imposed by requiring paths to stay on the grid graph. See also Mitchell et al. [MPK] and Thorpe [Th1,Th2]. These grid-based planners do not, however, solve the exact Euclidean shortest path problem.

For maps given in the form of a list of polygonal obstacles, there are two basic approaches to finding shortest paths. A survey of these methods is given in Mitchell and Papadimitriou [MP1] or Mitchell [Mi2].

The first approach is to build a *visibility graph* of the obstacle space, whose nodes are the vertices of the obstacles and whose edges connect pairs of vertices for which the segment between them does not cross the interior of any obstacle. The resulting graph can be searched using Dijkstra’s algorithm or A^* , where the heuristic function may again be taken to be the straight-line distance from a point to the goal. The time required for searching the graph is $O(K + n \log n)$, where K is the number of edges in the visibility graph. The result of the search is actually more than just the shortest path from START to GOAL; we actually get a *shortest path tree* on the set of obstacle vertices which gives a shortest path from the START (or the GOAL) to every other obstacle vertex.

The bottleneck in the worst-case complexity for the visibility graph approach is in the construction of the graph. We need to determine for every pair of vertices whether or not the vertices are visible to each other. There are $O(n^2)$ such pairs. A trivial $O(n^3)$ algorithm simply checks every pair of vertices against every obstacle edge. With a little more care, $O(n^2 \log n)$ suffices to compute the visibility graph (Lee [Le], Sharir and Schorr [SS], or Mitchell [Mi2]). The basic idea is to sort (angularly) the set of all vertices about each vertex and then to do an angular sweep about each vertex, keeping track of the closest obstacle boundary. Using a trick from duality theory that maps the vertices to lines, Lee and Ching [LC] showed that the n sorts can be done in total time $O(n^2)$ rather than $O(n^2 \log n)$. This fact allowed Welzl [We] and Asano et al. [AAGHI] to achieve an $O(n^2)$ algorithm for constructing visibility graphs. This running time is not likely to be improved for the worst case since there are visibility graphs of quadratic size.

However, the most recent development in the analysis of visibility graphs is the algorithm of Ghosh and Mount [GM] which computes the visibility graph of a set of n disjoint line segments in time $O(K + n \log n)$, where K is the number of edges in the visibility graph. This *output-sensitive* algorithm takes advantage of the fact that K is not always of size $\Theta(n^2)$, but rather can be as small as $\Theta(n)$. So, when the visibility graph is sparse (say, $O(n \log n)$ edges), the algorithm is very fast, building the visibility graph in time $O(n \log n)$. Thus, when $K = O(n \log n)$, Dijkstra’s algorithm (or A^*) requires worst-case complexity $O(K + n \log n) = O(n \log n)$, so the shortest path can be found in time $O(n \log n)$.

A second technique for finding shortest paths is that of building a *shortest path map*. A shortest path map is a subdivision of the plane into regions each of which is the locus of all goal points whose shortest paths from the START have the same topology (i.e., pass through the same sequence of obstacle vertices).

With such a subdivision, the shortest path from START to GOAL is obtained simply by locating the GOAL point in the subdivision (which takes time $O(\log n)$) and then backtracing an optimal path. See Lee and Preparata [LP] or Mitchell [Mi2] for precise definitions and examples. Reif and Storer [RS] have given an algorithm to compute the shortest path map in time $O(kn + n \log n)$, where k is the number of obstacles. Note that this is an improvement over $O(n^2)$ when the number of obstacles is small in comparison with the number of obstacle vertices. Their algorithm simulates the propagation of a “wavefront” from the START, in much the same way that expansion operates in the Dijkstra or A^* algorithms. Because of its similarity to the discrete Dijkstra algorithm, this technique has come to be called the “continuous Dijkstra” paradigm.

The continuous Dijkstra paradigm, in the form of a plane sweep algorithm, has led to $O(n \log n)$ algorithms in a few special cases (Lee and Preparata [LP], Mitchell [Mi2], Sharir and Schorr [SS]). Basically, if the obstacles are known to have disjoint projections onto some line (e.g., if the obstacles are vertical line segments), then the shortest path will always be monotone with respect to that line (except, perhaps, at the ends of the path). This monotonicity allows one to use plane sweep algorithms for the iterative construction of the shortest path map.

In the simple case of finding shortest paths required to stay inside a single simple polygonal room with n sides (without other obstacles), the shortest path can be found in linear ($O(n)$) time once a triangulation of the polygon is known (Lee and Preparata [LP], Guibas et al. [GHLST]). (Triangulation of a simple polygon is now known to require $o(n \log n)$ time, as Tarjan and Van Wyck [TV] have provided an $O(n \log \log n)$ algorithm. It is suspected that triangulation can be done in linear time, but this is another outstanding open problem in computational geometry.) Furthermore, a triangulated simple polygon can be processed in linear time so that in time $O(\log n)$ one can find the length of the shortest path from any given START to any given GOAL (Guibas and Hershberger [GH]). Furthermore, in time proportional to the number of bends in the optimal path, one can output the actual path from START to GOAL.

Lower bounds are known for the shortest obstacle-free path problem. In the worst case, computing the shortest path among k obstacle with a total of n vertices requires $\Omega(n + k \log k)$ time. One can use a shortest path algorithm to find convex hulls, which in turn can be used to sort integers, for which the lower bound is well known.

The big open problem in this field is to determine whether or not the shortest path among n disjoint line segments can be computed in optimal time $\Theta(n \log n)$, and more generally if shortest paths among k simple polygons can be found in time $\Theta(n + k \log k)$. We should note that for some metrics other than Euclidean, such as the L_1 (“Manhattan”) metric and “fixed orientation” metrics (Widmayer et al. [WWW]), optimal or near-optimal shortest path algorithms are known (see Clarkson et al. [CKV], de Rezende et al. [DLW], and Mitchell [Mi3]).

3. Shortest Paths on a Surface

A natural generalization of the obstacle-avoidance problem in the plane is to consider the case of finding

shortest paths for a point which is constrained to move on a nonplanar (possibly nonconvex) surface. To see that this is a generalization of the above problem, consider the surface which results from making each obstacle into a very tall cylinder whose base is in the plane of motion. Then, one can show that the shortest path on such a (nonconvex) surface will stay in the plane and follow the shortest obstacle-avoidance path. Note that the problem we wish to solve is a calculus of variations problem: Find the shortest *geodesic* path between two points on a surface. Since the surface is modeled in a discrete manner, we refer to this problem as the *Discrete Geodesic Problem*.

Here, we assume that our vehicle is not able to fly, but is able to travel over all types of terrain with uniform efficiency. We can go up hills, down hills, and even perhaps on the underside of an overhang. (Our tires are “magnetic” and the ground has a very high iron content.) Such is the problem faced by the crawling insect in Figure 2. He desires to get to the cheese in the shortest possible path, but he does not have wings to be able to fly there. Along what route should he crawl?




Figure 2. The problem of navigating on a surface.

First, let us mention something about the representation of the problem. What is the map in this case? It is simply some representation of a surface in three dimensions. Formally, the surface might be represented as a polyhedral surface, giving faces, edges, and vertices, along with all of the usual adjacency relationships that allow one to give the faces on each side of any edge and to “walk around” a vertex while enumerating adjacent faces. (The winged-edge data structure of Baumgart [Bau] or the similar quad-edge data structure of Guibas and Stolfi [GS] would work nicely here.) Another possible surface representation that is often used in practice is that of contour lines. We might be given a set of iso-elevation curves, as is usually the case with geological maps of terrain. A third common representation is that of an *elevation*

array. In this case, we are simply given a two-dimensional array of numbers which represent the altitude at each grid point. Digital terrain data bases of the form compiled by the Defense Mapping Agency fall into this category. (Typically, pixels are of size 5, 12.5, or 100 meters.) These are just some of the methods of specifying a discrete approximation of a surface. Note that the last two representations require that the surface be the image of a single-valued function of the plane (i.e., there are no “overhangs”), an assumption that generally holds for most outdoor terrain.

The case of a surface represented with an elevation array can be solved easily if one makes the approximation that paths are constrained to follow the corresponding grid graph. We just connect each pixel to each of its eight neighbors, assigning a length measured by the three-dimensional Euclidean distance, and then search the graph for a shortest path using Dijkstra or an A^* algorithm. This requires $O(n \log n)$ time in the worst case, where n is the number of array elements (pixels).

Research in the field of algorithmic motion planning has favored the representation of a surface in terms of a polyhedron. The general problem of finding shortest paths for a point in three dimensions amidst polyhedral obstacles has been of interest for some time (Sharir and Schorr [SS], Reif and Storer [RS], Franklin and Akman [FA], Akman [Ak1, Ak2], Papadimitriou [Pa], Clarkson [Cl]). The general problem seems very hard to solve, as Canny and Reif [CR] have shown the problem to be NP-hard, and the best known exact algorithm is the singly-exponential algorithm of Reif and Storer [RS]. Papadimitriou [Pa] and Clarkson [Cl] have solved the problem approximately with fully polynomial approximation schemes (meaning that the running time is polynomial in n , the number of bits, K , to represent the problem instance, and $1/\epsilon$). The complexity of their algorithms is $O(n^3 K^2/\epsilon)$. The special case, though, of our bug crawling on the surface of a single polyhedron is interesting because it admits a polynomial-time algorithm which is exact.

Franklin and Akman [FA] and Akman [Ak1,Ak2] have given exponential (worst-case) algorithms for the shortest path problem on a surface, and have shown that their algorithms perform reasonably well in practice. The first polynomial-time algorithm was an $O(n^3 \log n)$ algorithm of Sharir and Schorr [SS] which worked only for the case of convex polyhedra (where n is the number of edges of the polyhedron). Most interesting terrain is not convex, though, since it will usually have many mountains and valleys. An algorithm which works for non-convex surfaces was devised by O’Rourke et al. [OSB], and its complexity is $O(n^5)$. Both of these running times were improved by the algorithm of Mitchell et al. [MMP], which runs in time $O(n^2 \log n)$. The algorithm actually constructs, for a fixed start point, a subdivision (a *Shortest Path Map*) of the surface such that after the $O(n^2 \log n)$ preprocessing required to build the subdivision, the distance to any particular destination can be found by locating the goal in the subdivision (a task which is well-known to take $O(\log n)$ time; see, for example, Kirkpatrick [Ki]). Listing of a path can be done in time proportional to the number of turns in the optimal path.

Let us make a few simple observations that were critical to all of the algorithms mentioned above. We can state a few known facts about how optimal paths must behave:

1). Optimal paths are simple (not self-intersecting), and they can pass through any one face at most once.

2). Optimal paths “unfold” into straight lines. This is an observation that any high school geometry student has made if he has tried to solve the problem of finding the shortest path between two points on the surface of a box: Simply “cut” the box and flatten it in such a way that the start and goal can be connected by a straight line segment that stays on the flattened cardboard. One such unfolding is guaranteed to produce a shortest path. An optimal path which passes through an edge of the polyhedron must obey the *local optimality criterion* that it unfolds to give a straight line. More generally, for any type of surface there is a similar statement of the principle of optimality, and this produces a characterization of the *geodesics* of that type of surface. (Recall that the geodesics of a sphere are arcs of great circles; the geodesics of a polyhedral surface are characterized by the above unfolding property.) Of course, the question that must be answered by an algorithm is which of the (possibly exponentially many) unfoldings yields a shortest path?

3). On a convex surface, optimal paths will not go through vertices (except perhaps at their endpoints), while on a non-convex surface, an optimal path may pass through a vertex only if the angle it makes in so doing is greater than π (see Mitchell et al. [MMP] for a more precise statement).

Obviously, any algorithm which attempts to solve a shortest path problem should exploit as much of the structure of the paths as possible. It is also sometimes possible to exploit special structure when designing good heuristics. For example, to solve a shortest path problem on a given surface, we could use the simple heuristic which connects the start with the goal along the path “as the crow flies” (that is, the intersection of the surface with the vertical plane through the start and goal), and then iteratively improves the path by applying the local optimality criterion of (1) above. This will result in a *locally* optimal path which will frequently be very good in practice (while it is certainly not guaranteed to be *globally* optimal). Perhaps a planner which combines some “intelligence” with this heuristic would produce reasonably good paths. For example, one might have knowledge that there is only one feasible way through a mountain range, say through a specific gorge, and we would then want to route our initial guess through this passage, so that the locally optimal path produced will be reasonable. Further experimental and theoretical research is needed on these problems.

4. Weighted Regions

We now consider a more general terrain navigation problem. Assume that our map represents a set of *regions* in the plane, each of which has an associated *weight*, or cost, α . The weight of a region specifies the “cost per unit distance” of a vehicle (considered to be a point) traveling in that region. Our objective is then to find a path from START to GOAL in the plane which minimizes total cost according to this weighted Euclidean metric. For example, the weights may represent the reciprocal of the maximum speed in each

region, in which case minimizing the weighted Euclidean length of a path is simply minimizing the length of time it takes the vehicle to execute the path.

This problem has been termed the *Weighted Region Problem*. Note that it is a generalization of the obstacle-avoidance problem in the plane, for the obstacle-avoidance problem is simply the weighted region problem in which the “weights” are either 1 or $+\infty$ depending on whether a region is “free space” or obstacle, respectively. Thus, an important theoretical reason for considering the weighted region problem is that it is a natural generalization of the well-studied obstacle-avoidance problem.

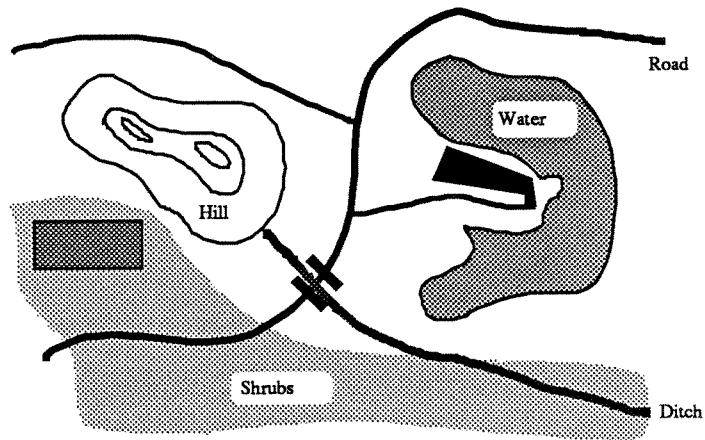


Figure 3. Map for terrain navigation.

An example of a typical map is illustrated in Figure 3. Basically, the ground surface is subdivided into uniform regions, with a label (or some traversability index) attached to each region which gives information about how fast one can move in that region (or how costly it is to do so). Presumably, the robot can move through different types of terrain at different speeds. Speeds may also depend on other factors such as time of day, precipitation, or the location of other vehicles; however, we assume that a given problem instance has fixed weights. In military applications, there may be regions which correspond to high threat risk, perhaps because the enemy has good visibility of you when you are in these regions. Costs can be assigned to traveling in these risky regions as well.

We will discuss two basic types of map representations: regular tessellations (e.g., grids of pixels, or quadtrees), and straight-line planar subdivisions.

Representing terrain in the form of a regular grid of pixels is natural and simple. Figure 4 shows the map of Figure 3 in digitized form, on a grid of size 26-by-40 (1040 pixels). Frequently, terrain data is given in the form of a set of arrays, with each array giving information about some aspect of the terrain (e.g., ground cover, land usage, hydrography data, man-made features, and traversability indices). This, for example, is the form in which the Defense Mapping Agency terrain data is supplied. Pixels are usually squares 5, 12.5, or 100 meters on a side. From this type of data, we can specify the map by taking the

superposition of all the data arrays to form one composite map (a “terrain array”), whose regions of equal attributes will have descriptions such as “brush-covered, drainage plain”, or “forested land, sandy soil, small boulders”. Each composite region may be assigned a weight representing the cost of motion in that type of terrain. Certain types of terrain features have priority over others. For example, if the “roads array” tells us there is a road surface occupying pixel (i, j) , while the ground-cover array describes the pixel as “grassland”, then the composite region attribute at (i, j) should say that there is a road there.




Figure 4. Digitized terrain map.

If the terrain map is given as a grid of weighted pixels, then a straightforward solution to the weighted region problem is to search the corresponding grid graph for shortest paths (Jones [Jo], Keirse and Mitchell [KM], Quek et al. [QFP]). This is simply a generalization of the approach described earlier for the case of digitized obstacle maps. The grid graph will have either 4- or 8-connectivity depending on whether or not we consider diagonal neighbors of pixels. Costs may be assigned in the natural way to arcs connecting adjacent pixels of the grid, taking into account the weights of the two pixels by, say, taking the cost of the arc between them to be their average weight. Then, Dijkstra’s algorithm or A^* may be used to compute minimum-cost paths in the grid graph. For a grid with n pixels, this results in an algorithm that runs in worst-case time $O(n \log n)$. This approach is particularly appealing in cases in which the data is given to us in the form of feature arrays.

One problem with the grid graph approach is that it may require extremely fine grids to capture the content of a relatively simple piece of terrain. Frequently, there may be just a couple of large uniform regions, perhaps with a road running through them, and the pixel representation of the map requires, say, a 512-by-512 array. The shortest path may indeed be obvious (such as a simple straight line segment), but the algorithm must still search and expand thousands of pixels during its execution (although hierarchical algorithms can avoid much of the expansion). Another problem with the grid graph approach is that it creates a *digitization bias* because of the metrication error imposed by confining movements to 4 or 8 orientations.

See Keirsey and Mitchell [KM] for a more detailed discussion of digitization bias and the various possible remedies for it. See Quek et al. [QFP] for a technique that uses grid graphs of multiple resolutions to design a hierarchical algorithm.

If, instead of being a regular tessellation, the map is modeled as a straight-line planar subdivision then the regions are just simple polygons. Figure 5 shows the map of Figure 3 drawn as a straight-line planar subdivision with 85 edges. An appropriate data structure would be either a winged-edge (Baumgart [Bau]) or a quad-edge (Guibas and Stolfi [GS]) data structure. Roads and other linear features (such as ditches, fences, streams, etc.) can be modeled as very skinny polygonal regions or approximated as sets of line segments. In the planar subdivision, then, we allow edges to be assigned a weight which may be different from the weights of the faces on either side of it. If an edge is a line segment of a road, then, presumably, its weight will be less than the weights on either side of it (assuming that it is cheaper to travel on roads than off roads). If an edge is a line segment of a ditch or a fence, then it may be assigned a fixed cost of crossing it. The fixed cost would be $+\infty$ if the edge cannot be crossed.




Figure 5. Polygonal subdivision map.

The representation of regions as polygonal patches and of roads as linear features is being used in the ongoing work on the DARPA Autonomous Land Vehicle project. One advantage of encoding the map information originally in the form of polygonal patches is that it can save greatly on the storage costs, while increasing the resolution of the data. The centerline of a road need not be specified only to the resolution of a pixel (say 12.5 meters), but can be specified to any desired degree of accuracy and can be specified to a higher resolution than other less critical features of the terrain.

In order to put the terrain array into the representation of a polygonal subdivision, we must fit polygonal boundaries about the sets of pixels with uniform attributes. At the highest resolution, each pixel can be considered to be a polygonal patch, namely, a square; however, this extreme will usually result in many more

regions than desired, and it does not reflect the fact that the original digital data was only approximate in the first place. Usually, we will also want to represent roads, fences, and power lines as linear features by fitting a piecewise-linear path to them.

Assume now that we are given a planar subdivision with weights. One approach to solving the weighted region problem is to build a “region graph” in which nodes correspond to regions and arcs correspond to boundaries between adjacent regions. Assume that the subdivision is fine enough that all regions are convex. (We could, for example, make the subdivision a triangulation by triangulating each of the simple polygonal regions; however, we may want to subdivide in such a way that each convex region is as close to being circular as possible.) Then we can think of placing a node at the “center” (e.g., center of mass) of each region. Two nodes are joined by an edge if the corresponding regions are adjacent. We then assign costs to arcs according to the weighted distance between adjacent nodes. (An alternative graph can be constructed by placing nodes at the midpoints of region edges and linking two nodes if the corresponding edges share a common region.) Note that a grid graph is a special case of a region graph in which all regions are equal-sized squares. Searching this graph for shortest paths yields a “region path” from the START to the GOAL, giving a sequence of regions through which a “good” path should pass. We could then do some post-processing (e.g., using Snell’s Law of refraction from optics, a local optimality criterion that we will mention below) to make the path locally optimal at region boundaries. The problem with this approach is that it can produce paths that are far from being optimal, for the optimal path need not have any relationship to the shortest region-path. In practice, however, it promises to be a very useful technique, particularly in cases in which the data is not very accurate or the assignment of weights is somewhat subjective. (Can one really say that grassland deserves a weight of 5 while brushland deserves a weight of 15? It seems likely that weights will be highly subjective.)

Thus, the region graph approach does not guarantee that the resulting path is close to being optimal. Another approach given by Mitchell and Papadimitriou [MP2] yields a polynomial-time algorithm for computing ϵ -optimal paths. The length of the produced path is guaranteed to be within a factor of $1 + \epsilon$ of the length of an optimal path. The complexity of the algorithm is $O(n^7L)$, where L is the number of bits needed to represent the problem data and ϵ . The algorithm is a generalization of the algorithm for the discrete geodesic problem given by Mitchell et al. [MMP], and it too uses the continuous Dijkstra paradigm.

Let us be more precise about the problem solved in [MP2]. We are given a (straight-line) planar subdivision, \mathcal{S} , specified by a set of faces (regions), edges, and vertices, with each edge occurring in two faces and two faces intersecting either at a common edge, a vertex, or not at all. We consider faces to be *closed* polygons (they include their boundaries) and edges to be *closed* line segments (they include their endpoints, which are vertices). We are also given an initial point, START. Without loss of generality, we assume that all faces are triangles and that s and t are vertices of the triangulation. Assume that \mathcal{S} has n edges (and hence $O(n)$ triangular faces and $O(n)$ vertices). Our complexity measures will be written in terms of n .

Each face f has associated with it a weight $\alpha_f \in [0, +\infty]$ which specifies the cost per unit distance

of traveling on face f . Similarly, each edge e has associated with it a weight $\alpha_e \in [0, +\infty]$. The weighted distance between any two points x and y on edge e is simply the product $\alpha_e|xy|$, where $|xy|$ is the usual Euclidean distance between x and y . Likewise, the weighted distance between any two points x and y on face f (but not both on the same edge of f) is the product $\alpha_f|xy|$. The weighted length of a path through the subdivision is then the sum of the weighted lengths of its subpaths through each face.

We are asked to find the minimal length (in the weighted sense) path from START to some goal point. The algorithm of [MP2] actually solves the *query* form of the weighted region problem: Build a structure which allows one to compute a shortest path (in the weighted Euclidean metric) from START to any query point t . It turns out that this single-source version of the problem is no harder (in worst-case asymptotic time) than the problem with a given fixed goal point.

As with the discrete geodesic problem, the algorithm for the weighted region problem exploits certain facts about the local behavior of optimal paths:

1). Shortest paths are piecewise-linear, bending only at points where they cross an edge or pass through a vertex. This follows from the fact that each region is assumed to have a *uniform* weight. (The case of nonuniform weighted regions gives rise to other types of curves and is an interesting area for further research.)

2). When an optimal path passes through an edge between regions, it does so while obeying Snell's Law of Refraction at that edge. This is the local optimality criterion which is analogous to the fact that shortest paths on the surface of a polyhedron unfold to be straight lines. It is a simple property of locally-optimal paths through weighted regions which has been observed by many researchers (Lyuaternik [Ly], Mitchell and Papadimitriou [MP1,MP2], Richbourg [Ri], Richbourg et al. [RRZ,RRZM], Rowe and Richbourg [RR], Smith [Sm]).

Snell's Law of Refraction *The path of a light ray passing through a boundary e between regions f and f' with indices of refraction α and β obeys the relationship that $\alpha \sin \theta = \beta \sin \phi$, where θ and ϕ are the angles of incidence and refraction (respectively).*

The *angle of incidence*, θ , is defined to be the counterclockwise angle between the incoming ray and the normal to the region boundary. The *angle of refraction*, ϕ , is defined as the angle between the outgoing ray and the normal. Refer to Figure 6.

The fact that light obeys Snell's Law comes from the fact that light seeks the path of minimum time (this is *Fermat's Principle*). The index of refraction for a region is proportional to the speed at which light can travel through that region. Hence, the shortest paths in our weighted region problem must also obey Snell's Law. This can be shown formally (Mitchell and Papadimitriou [MP2], Mitchell [Mi2]).

3). The angle of incidence of an optimal path at a boundary that it crosses will not be greater (in absolute value) than the *critical angle* defined at that boundary. This fact comes from examining the relationship specified in Snell's Law and making sure it is well-defined. Without loss of generality, assume




Figure 6. Light ray crossing a boundary.

that $\alpha_f > \alpha_{f'} > 0$. Then, we must assure that

$$-1 \leq \sin \alpha_{f'} = \frac{\alpha}{\alpha_{f'}} \sin \theta \leq 1.$$

The angle, $\theta_c(e) = \theta_c(f, f')$, at which

$$\frac{\alpha_f}{\alpha_{f'}} \sin[\theta_c(f, f')] = 1$$

is called the *critical angle* defined by edge e . If $\alpha_f > 0$ and $\alpha_{f'} \geq 0$, then the critical angle is defined and is given by $\theta_c(f, f') = \sin^{-1}(\alpha_{f'}/\alpha_f)$, provided that $\alpha_f \geq \alpha_{f'}$. A ray of light which strikes e at the critical angle will (theoretically) travel along e rather than enter into the interior of f' . We can think of it as if the light ray wishes to be “just inside” the “cheap” region on the other side of e .

4). Optimal paths will not reflect from a boundary, except when they are “critically reflected” as shown in Figure 7: A path is incident (from a face f of weight β) on edge e (of weight $\gamma < \min\{\alpha, \beta\}$) at the critical angle $\phi = \theta_c$ at some point $y \in \text{int}(e)$, it then travels along edge e for some positive distance, and then exits edge e back into face f at some point $y' \in \text{int}(e)$, leaving the edge at an angle $\phi = \theta_c$. We then say that the path is *critically reflected* by edge e and that segment $\overline{yy'}$ is a *critical segment* of p along e . The phenomenon of critical reflection is not just an unlikely event of degeneracy. It is commonly the case that the shortest path between two points consists of a single critical reflection along an edge.

5). If $\alpha_e > \min\{\alpha_f, \alpha_{f'}\}$, then an optimal path will never travel along edge e , since it would be better to travel “just inside” either region f or region f' (whichever is cheaper).

6). An optimal path can cross a “road” (that is, an edge whose cost is less than that of the regions on either side of it) in one of two ways: either by obeying Snell’s Law while passing through the boundary, or by “hitching a ride” along the edge, as illustrated in Figure 8. The path hits edge e at the *incoming* critical




Figure 7. An optimal path which is critically reflected.




Figure 8. Shortest path "hitching a ride" on a road.

angle $\phi = \theta_c(f, e) = \sin^{-1}(\alpha_e/\alpha_f)$, then travels along edge e (of weight γ) for some distance, then leaves edge e into f' at the *outgoing* critical angle $\theta = \theta_c(f', e) = \sin^{-1}(\alpha_e/\alpha_{f'})$.

7). On an optimal path, between any critical point of exit and the next critical point of entry, there must be a vertex of \mathcal{S} . This fact turns out to be very important in the proofs of the complexity of the algorithm of Mitchell and Papadimitriou [MP2].

The local optimality criteria outlined above are sufficiently strong that they uniquely specify a locally optimal path which passes through a given sequence of edges. (This follows from the global convexity of the objective function when written in terms of the coordinates at which the path crosses each edge.) The problem remains, however, of how to select the proper sequence edges for the optimal path.

The continuous Dijkstra algorithm of Mitchell and Papadimitriou [MP2] simulates propagation of a wavefront from START. Each time the wavefront first hits a region boundary, passes through a vertex, or otherwise changes its description (that is, the list of regions through which it passes) we record the effect of the event in an appropriate data structure. When the wave has propagated throughout the surface, we have completed the construction of a shortest path map. For the application to the weighted region problem, it can be shown that the worst-case number of event points is $O(n^3)$, thereby establishing a polynomial time bound.

Processing each event involves calling certain functions. One such function locates a point in an “access channel” between representative optimal paths that have already been established. This is easy and can be done by binary search in $O(\log n)$ time. However, we must also do $O(n^3)$ operations of the form “Find the refraction path from a given point r that passes through a given sequence of edges (while obeying Snell’s Law) to hit a given point x ”. This is not an easy problem to solve exactly. It is handled by a numerical routine that does binary search to find a path whose length is at most $(1 + \epsilon)$ times the length of the optimal path, where $\epsilon > 0$ is an arbitrary error tolerance value that is given in the problem specification. The search requires time which is polynomial in n and L , the number of bits necessary to represent the problem instance (namely, $\log(nNW/\epsilon w)$, where N is the largest integer value of any vertex coordinate, W is the largest finite weight, and w is the smallest non-zero weight). A crude calculation shows that the time to run the numerical search is bounded by $O(n^4L)$.

It would be great if we could calculate exactly the refraction path from point r to point x , but this problem seems to be hard. The complication is that when we write down the equations that represent the local optimality criterion (Snell’s Law) at each edge, and then try to solve for the points where the path crosses edges, we get k quartic equations in k unknowns (where k is the length of the edge sequence through which the refraction path is known to pass). Elimination among these equations yields a very high degree polynomial (of degree exponential in k). Alternatively, we could apply the cylindrical decomposition technique of Collins [Co] to arrive at a doubly exponential time procedure to determine whether or not a given rational path length is achievable. This would also provide us with a technique of comparing the lengths of paths through two different edge sequences, and thus would solve the entire problem precisely in doubly exponential time (outputting the sequence of edges and vertices along the optimal path). The same technique led Sharir and Schorr [SS] to a doubly exponential time solution to the three-dimensional shortest path problem.

In practice, a very straightforward numerical approach to solving the search problem can be used. For example, one technique begins with a path from r to x that connects the edges in the sequence along their midpoints. We then iteratively shorten the path by applying the local optimality criterion to adjacent segments of the path. We simply pick an edge at which Snell’s Law is violated, and then let the crossing point “slide” along the edge until Snell’s Law is obeyed. (This is a coordinate descent method for searching for the minimum of the convex function which describes the path length.) Each iteration can be computed

in constant time and results in a strict decrease in the weighted path length. Hence, this procedure will converge to the locally optimal path from r to x , which is also the retraction path from r to x . (This type of numerical procedure is also applicable to the three-dimensional shortest path problem, in which the local optimality criterion is that paths unfold to be straight.) The problem with this approach is that we do not have good bounds on the rate of convergence or on the number of iterations necessary to guarantee that the solution is close (say, within $\epsilon\%$) to optimal.

In practice, however, the convergence rate has been observed to be “very fast”. In particular, in some recent implementation work of Karel Zikan [Zi], he has observed that a version of the above-mentioned coordinate descent method tends to take only about 5-10 iterations to get within 1 percent of optimality for problems with up to about 20 edges. By employing certain “overrelaxation” methods, he is able to improve the convergence rate, cutting approximately in half the number of iterations needed.

5. Special Cases of the Weighted Region Problem

There are a few special cases of the weighted region problem which admit alternative polynomial-time algorithms and are of interest in terrain navigation. We have already mentioned the special case in which all weights are 1 or $+\infty$. This leads to the usual obstacle-avoidance problem in the plane.

Consider now the special case in which all the weights on faces and edges are 0, 1, or $+\infty$. Infinite-cost regions are obstacles, cost one regions are simply “free space” (through which we can move at some fixed speed), and cost zero regions are “freebies” (through which motion costs nothing; we can move at infinite speed). If there are no zero-cost regions, then the problem is simply the usual obstacle avoidance shortest path problem in the plane. When there are zero-cost regions, they behave like “islands” between which one may want to “hop” to get from the source to the destination.

As a motivation for this special case, let us first note that it may be a reasonable approximation to the case in which the terrain is divided into three categories: obstacles (or very costly regions), very low cost regions, and regions whose costs are in between, but similar. Such would be the case for a runner (who is not a very good swimmer) trying to find a route from one point (say, on a flat island) to another point (say, in the water). There are many islands around. Some are flat and can be crossed easily on foot. Others have huge cliffs surrounding them which make them impassable. Swimming through the water is possible, but the runner would much prefer to be running across a flat island. So how can he get from one point to another in the shortest time? His path will “hop” from one flat island to another, swimming in between, and avoiding the mountainous islands.

An application of this problem which has arisen in practice is the *maximum concealment* problem. There are regions of the plane which are “visible” to an enemy threat, and these are hopefully to be avoided. There are also obstacles, which *must* be avoided. Then there are regions (e.g., between mountain ranges, or behind rocks) which are hidden from the enemy’s view, and are hence extremely cheap in comparison with the visible regions. See Figure 9 for an example in which there is a single enemy observer whose view is

obstructed by tall obstacles, but who can see over the bodies of water. The tall obstacles cast “shadows”, which are regions of essentially free travel. The problem is to find an obstacle-avoiding path between two points which maximizes the concealment (minimizing the length of time in the exposed regions). In the context of the weighted region problem, the model is that water and mountain ranges are infinite-weight regions, shadowed regions behind mountains are zero-cost regions (for purposes of concealment), and all other terrain has weight 1. (It makes sense that the cost of being seen by a threat should depend on the distance from the threat. Models that make this assumption are discussed in Mitchell [Mi6], where polynomial-time algorithms are given for some special cases.)




Figure 9. Maximal concealment from a single enemy observer.

Other applications also use the $\{0, 1, +\infty\}$ special case of the weighted region problem. It can be used to find the shortest path from one *region* to another *region* in the presence of obstacles (by letting the source and destination regions be zero-weight regions), as was independently considered by Asano et al. [AAI]. It can also be used to find *lexicographically* shortest paths through weighted regions: minimize the length of path in the most expensive region, then, subject to this being minimal, minimize the length in the next most costly region, etc. Lexicographically shortest paths may be good initial guesses for heuristic algorithms that try to find shortest paths through weighted regions. Also, there is an intimate connection between shortest path problems and *maximum flow* problems, and the algorithms for maximum flow need to find shortest paths through $\{0, 1, +\infty\}$ weighted regions. These and other applications are discussed in Mitchell [Mi4, Mi5].

An interesting application discussed in Gewali et al. [GMMN] is that of the “least-risk watchman route problem”. The watchman problem asks one to find the shortest path that a watchman would have to make from his current location in order to be able to “see” all of a given space (e.g., a polygonal room, or a polygonal room with “holes”). Chin and Ntafos [CN1,CN2] have shown that the optimum watchman route problem is NP-complete in a polygonal room with holes, but that there is a polynomial-time algorithm for polygonal rooms without holes. The *least-risk watchman route problem* asks one to find the path of minimum exposure to threats which sees the entire space. Then, for rectilinear polygons with n sides, Gewali et al. [GMMN] give an $O(k^2n^3)$ algorithm for the case of k threats.

The algorithm described by Mitchell and Papadimitriou [MP2] for the weighted region problem does in fact solve the $\{0, 1, +\infty\}$ special case in polynomial time. But the special case has structure which allows an alternative approach which is faster and easier to implement. The approach in Mitchell [Mi4] and Gewali et al. [GMMN] to this case is to build a special kind of “extended visibility graph”, VG^* , after shrinking the zero-cost regions to single nodes. The graph VG^* can be constructed in polynomial time (in fact, quadratic time), and it can then be searched for shortest paths in the usual way using Dijkstra or A^* . The result is an exact polynomial-time algorithm. The basic idea is to exploit the local optimality criteria, which allow the problem to be discretized by limiting our search to the graph VG^* .

1). The behavior of optimal paths with respect to the obstacles is the same as it was in the obstacle-avoidance problem (see Mitchell [Mi2]). Namely, the path behaves around obstacles as if it were a “taut string”, being locally tangent to obstacles and forming a convex angle when it comes in contact with an obstacle vertex.

2). An optimal path will either enter a zero-cost region at a vertex or at a normal to the interior of an edge bounding the region. If it enters at a vertex, then it must make an angle greater than $\pi/2$ with the edges incident to the vertex.

Thus, the extended visibility graph has edges of two types: those that join two vertices (of obstacles or zero-cost regions), and those that join a vertex (of either an obstacle or a zero-cost region) to an edge of a zero-cost region (in such a way as to be perpendicular to the edge). This immediately implies a quadratic bound on the size of the graph. The exact details of the construction are given in either of the above-mentioned papers ([Mi4] or [GMMN]). The result is an $O(n^2)$ algorithm for shortest paths.

Another special case of interest in terrain navigation is that in which all of the terrain has the same (finite) weight, but there exists a (piecewise-linear) network of roads or other linear features (of possibly different sizes, which admit different speeds of traversal). This is the case illustrated in Figure 10, and is simply the weighted region problem in which all face weights, α_f , are identical, and edge weights, α_e , are arbitrary. Then, we apply the following local optimality criterion:

3). Shortest paths will enter a road segment either at one of its endpoints, or by hitting at the critical angle of incidence defined by the ratio of the edge’s weight to that of the surrounding terrain.




Figure 10. A map consisting of just roads.

This observation again leads to a discrete graph (a *critical graph*) which can be searched for shortest paths. Mitchell [Mi4] and Gewali et al. [GMMN] show how to construct this critical graph and search it in time $O(n^2 \log n)$, where n is the number of road edges. See also Rowe [Ro].

A generalization of this technique is possible by putting together all of the above-mentioned local optimality criteria. Assume now that we have a weighted region problem in which the weights on faces are in the set $\{0, 1, +\infty\}$, and the weights on edges are arbitrary ($\in [0, +\infty]$). Additionally, there can be a fixed cost, ξ_e , associated with crossing an edge e . The model is that there are obstacles, zero-cost regions, and background regions (all of the same weight). Additionally, there are roads, ditches, fences, etc. of arbitrary weights. Then, the algorithm given in Mitchell [Mi4] or Gewali et al. [GMMN] finds shortest paths in time $O(n^2 \log n)$, again by building a critical graph.

6. Generalizations and Extensions

Several generalizations to the discrete geodesic and weighted region problems are possible. Some are solved by minor changes to the algorithms of Mitchell et al. [MMP] and Mitchell and Papadimitriou [MP2], while others are open research problems.

1). First, one can generalize to the case of *multiple source points*. This means that instead of building a shortest path map for one point, START, we consider several given point sites (called “source points”), and we compute a subdivision of the plane or the surface which tells us for each point t which source point is closest to it and how to backtrace a shortest path to the closest source point. The resulting structure is then a classic *Voronoi diagram* according to the metric of interest (either shortest Euclidean paths on a surface or weighted Euclidean lengths through a weighted map).

As a further generalization along these lines, we could allow the source and goal to be *regions* rather than single points. This generality is actually already incorporated in the model of Mitchell and Papadim-

itriou [MP2] by allowing there to be zero-cost regions. One simply makes the starting region have zero weight, and places the starting point anywhere inside; a similar accomodation is made for the goal region. The applicability of this problem is apparent: Goals are frequently specified in the form “Get me to the town square,” or “Get me to Highway 95.” Any point within the square or on the highway will suffice as a goal.

2). We can generalize the cost structure from being that of a *uniform* cost per distance in a region to that of allowing a cost function (such as a linear function) in each region. The function should be specified by some fixed number of parameters, and should allow computation of geodesics. In the uniform cost case, the geodesics within a given region are simply straight lines. In more general cases, geodesics would have to be computed by techniques of calculus of variations (Lyusternik [Ly]). The algorithm must then be modified to apply the local optimality criterion (Snell’s law) to curved arcs at the boundaries of regions. If an optimal path passes through the point y interior to edge e , then the tangent lines at y to the geodesic curves in f and in f' must meet at y such that they obey Snell’s law. Using this local optimality criterion, then, it should be possible to run an analogous algorithm to that of Mitchell and Papadimitriou [MP2], this time piecing together curves from region to region. The details of this generalization need to be examined further.

3). We could allow boundaries between regions to be curved (e.g., splinegons; see Souvaine [So]). Then, for the weighted region problem, Snell’s law at a point y interior to a boundary curve must be applied as if the boundary at y is the straight line tangent to the boundary at y . Again, more research is needed on this problem.

4). The weighted region problem can be generalized to the case of a weighted polyhedral surface. Then the local optimality criterion becomes Snell’s law applied to the “unfolded problem” at each boundary. That is, for an edge $e = f \cap f'$, we must “unfold” f' about e so that f and f' are made coplanar, and then we can apply Snell’s law to points interior to e in this rotated coordinate frame.

5). We could allow the traversal of a region to be directionally dependent. This comment applies both to the (unweighted) discrete geodesic problem and to the weighted region problem. For example, it makes sense that it should be more costly to go “up” a hill than to go down it. We are currently investigating this case.

6). An important open problem is that of incorporating other types of vehicle constraints into the shortest path problem. For example, we may want a minimum time path for a vehicle which has a bounded amount of power (implying bounded acceleration). Or it may be that the vehicle will tip over if it is on too steep of a slope. Or the vehicle may have fuel consumption constraints. Solving the exact shortest path problem in the presence of other such constraints is an important area for continuing research.

7. Conclusion

We have examined some of the models of shortest path problems which are relevant to route planning for

an autonomous vehicle. We have discussed problems which admit exact solutions, and have discussed the currently best known algorithms. Many important issues remain to be addressed.

Obviously, there remains the question of reducing the worst-case running times of the algorithms discussed. In particular, it may be possible to improve the $O(n^7L)$ bound on the weighted region problem. It is a challenging problem to determine the *expected* running times of shortest path algorithms.

Problems involving the design of a good terrain map need to be studied more. We have discussed regular tessellations and polygonal subdivisions as representations for map data, and we have mentioned some of the benefits of each. What is probably needed is some kind of hybrid representation which can take advantage of the strengths of each, while also providing a hierarchical representation of the terrain.

If we were to compose a “wish list” of features for a map representation, we would most certainly want to include the ability to handle the following.

Discovery of new information. Ideally, as the vehicle moves about through terrain, it will update its map with new information it “learns”. It may, for example, discover that road construction is causing congestion along some stretch of highway or that a certain bridge it thought was there has been closed. It is therefore important to have a representation of a map which allows updating and annotations.

This also brings up the issue of interfacing the map-based path planner with some *sensor*-based planner that accumulates local information during the execution of a plan and must then modify the path accordingly. We have not addressed here the issue of planning routes in unknown or uncertain terrains, but we can refer the reader to a few of the myriad of papers on the subject: Goldstein et al. [GPSW], Iyengar et al. [IJRW], Jorgensen et al. [JHW], Kauffman [Ka], Kuan [Ku], Kuan et al. [KBZD], Kuan et al. [KZB], Lumelsky and Stepanov [LS], Meng [Men1, Men2], Metea and Tsai [MT], Mitchell [Mi1], Mitchell et al. [MPK], Moravec [Mo], Oommen et al. [OIRK], Rao et al. [RIJW1, RIJW2], Rueb and Wong [RW], Stern [St], Sutherland [Su], and Thorpe [Th2].

Temporal nature of information. Annotations might carry with them some estimated *duration* which bounds the interval of time for which the information is valid. For example, road construction is usually known to affect traffic flow only during certain times of the day, and it remains in effect only until it achieves completion.

Of course, uncertainty assures that the planner will never be able to guarantee that the routes it selects in advance are optimal. All types of map information are subject to some uncertainty. For example, roads can be closed or become impassable, open terrain can become flooded or reforested, boundaries of lakes and rivers certainly change with the seasons and with the rainfall pattern. A truly “intelligent” planner should have information about the kinds of circumstances which potentially affect different types of terrain and be able to reason about it. If it has knowledge, for example, that there was recently a greater than average rainfall, it should adjust the probabilities it places on the traversability of a low-water bridge.

Many open questions remain. How can these issues be addressed within the framework of exact or

approximate algorithms? What are the appropriate mathematical models of the path-planning process? How can our knowledge of algorithmic path planners be translated into useful heuristics for complex terrain problems?

8. Acknowledgment

This research was partially supported by a generous grant from the Hughes Aircraft Company. Part of this research was conducted while J. Mitchell was supported by a Howard Hughes Doctoral Fellowship at Stanford University and affiliated with the Hughes Artificial Intelligence Center in Calabasas, California. This research was also partially supported by a grant from the National Science Foundation (IRI-8710858).

9. Bibliography

- [Ak1] V. Akman, "Shortest Paths Avoiding Polyhedral Obstacles In 3-Dimensional Euclidean Space", PhD Thesis, Computer and Systems Engineering, Rensselaer Polytechnic Institute, June, 1985.
- [Ak2] V. Akman, "Unobstructed Shortest Paths in Polyhedral Environments", *Lecture Notes in Computer Science*, Vol. 251, Edited by G. Goos and J. Hartmanis, Springer-Verlag, 1987.
- [AAGHI] T. Asano, T. Asano, L. Guibas, J. Hershberger, and H. Imai, "Visibility of Disjoint Polygons", *Algorithmica*, Vol. 1, (1986), pp. 49-63.
- [AAI] T. Asano, T. Asano, and H. Imai, "Shortest Paths Between Two Simple Polygons", *Information Processing Letters*, **24**, 1987, pp. 285-288.
- [Bak] B. Baker, "Shortest Paths With Unit Clearance Among Polygonal Obstacles", *SIAM Conference on Geometric Modeling and Robotics*, Albany, NY, July 15-19, 1985.
- [Bau] B.G. Baumgart, "A Polyhedron Representation for Computer Vision", *AFIPS Conference Proceedings*, Vol. 44, 1975 National Computer Conference, pp. 589-596.
- [Br] R.A. Brooks, "Solving the Find-Path Problem by Good Representation of Free Space", *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-13**(3), pp. 190-197, March/April 1983.
- [BL] R.A. Brooks and T. Lozano-Pérez, "A Subdivision Algorithm in the Configuration Space for Find Path with Rotation", *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-15**(2), pp. 224-233, March/April 1985.
- [Ca] J.F. Canny, "The Complexity of Robot Motion Planning", Ph.D. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, May 1987.
- [CR] J. Canny and J. Reif, "New Lower Bound Techniques for Robot Motion Planning Problems", *Proc. 28th Annual IEEE Foundations of Computer Science*, pp. 49-60, Oct. 1987.
- [Ch] L.P. Chew, "Planning the Shortest Path for a Disk in $O(n^2 \log n)$ Time", *Proc. First Annual ACM Conference on Computational Geometry*, Baltimore, MD, pp. 214-220, 1985.
- [CN1] W.P. Chin and S. Ntafos, "Optimum Watchman Routes", *Proc. 2nd ACM Symposium on Computational Geometry*, pp. 24-33, 1986. (Submitted for publication.)

- [CN2] W.P. Chin and S. Ntafos, "Watchman Routes in Simple Polygons", Technical Report, Computer Science, University of Texas at Dallas, 1987.
- [Cl] K. Clarkson, "Approximation Algorithms for Shortest Path Motion Planning", *Proc. 19th Annual ACM Symposium on Theory of Computing*, New York City, May 25-27, 1987, pp. 56-65.
- [CKV] K. Clarkson, S. Kapoor, and P. Vaidya, "Rectilinear Shortest Paths through Polygonal Obstacles in $O(n \log^2 n)$ Time", *Proc. Third Annual ACM Conference on Computational Geometry*, Waterloo, Ontario, 1987, pp. 251-257.
- [Co] G.E. Collins, "Quantifier Elimination for Real Closed Fields by Cylindric Algebraic Decomposition", *Proc. Second GI Conference on Automata Theory and Formal Languages*, Springer-Verlag LNCS 35, Berlin, 134-183, 1975.
- [DLW] P.J. de Rezende, D.T. Lee, and Y.F. Wu, "Rectilinear Shortest Paths With Rectangular Barriers", *First ACM Conference on Computational Geometry*, June 1985, pp. 204-213.
- [Di] Dijkstra "A Note On Two Problems in Connection With Graphs", *Numerische Mathematik*, 1 (1959), pp. 269-271.
- [Ed] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, Germany, 1987.
- [FA] W. R. Franklin and V. Akman, "Shortest Paths Between Source and Goal Points Located On/Around a Convex Polyhedron", *22nd Allerton Conference on Communication, Control and Computing*, pp. 103-112, 1984.
- [FT] M. Fredman and R. Tarjan, "Fibonacci Heaps and Their Uses in Improved Network Optimization Algorithms", *Proc. 25th Annual IEEE Symposium on Foundations of Computer Science*, pp. 338-346, 1984.
- [GMMN] L. Gewali, A. Meng, J.S.B. Mitchell, and S. Ntafos, "Path Planning in 0/1 Weighted Regions With Applications", Extended Abstract, Department of Operations Research, Cornell University, 1987.
- [GM] S.K. Ghosh and D.M. Mount, "An Output Sensitive Algorithm for Computing Visibility Graphs", *Proc. 28th Annual IEEE Symposium on Foundations of Computer Science*, 1985.
- [GPSW] M. Goldstein, F.G. Pin, G. de Saussure, and C.R. Weisbin, "3-D World Modeling With Updating Capability Based on Combinatorial Geometry", Technical Report CESAR-87/02, Oak Ridge National Laboratory, 1987.
- [GHLST] L. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. Tarjan, "Linear Time Algorithms for Visibility and Shortest Path Problems Inside Simple Polygons", *Proc. Second Annual ACM Conference on Computational Geometry*, Yorktown Heights, NY, 1986, pp. 1-13.
- [GH] L.J. Guibas and J. Hershberger, "Optimal Shortest Path Queries in a Simple Polygon", *Proc. Third Annual ACM Conference on Computational Geometry*, Waterloo, Ontario, 1987, pp. 50-63.
- [GS] L. Guibas and J. Stolfi, "Primitives for the Manipulation of General Subdivisions and the Computation of Voronoi Diagrams", Technical Report, DEC Systems Research Center, 1985.
- [IJRW] S.S. Iyengar, C.C. Jorgensen, S.V.N. Rao, and Weisbin, "Robot Navigation Algorithms Using Learned

- Spatial Graphs”, *Robotica*, 4, (1986), pp. 93-100.
- [Jo] S.T. Jones, “Solving Problems Involving Variable Terrain. Part 1: A General Algorithm”, *Byte*, Vol. 5, No. 2, February, 1980.
- [JHW] C. Jorgensen, W. Hamel, and C. Weisbin, “Autonomous Robot Navigation”, *BYTE*, Vol. 11, No. 1, January 1986, pp. 223-235.
- [KD] S. Kambhampati and L.S. Davis, “Multiresolution Path Planning for Mobile Robots”, *IEEE Journal of Robotics and Automation*, **RA-2**, No. 3, September 1986, pp. 135-145.
- [Ka] S. Kauffman, “An Algorithmic Approach to Intelligent Robot Mobility”, *Robotics Age*, Vol. 5, No. 3, May/June 1983, pp. 38-47.
- [KM] D.M. Keirseay and J.S.B. Mitchell, “Planning Strategic Paths Through Variable Terrain Data”, *Proc. SPIE Applications of Artificial Intelligence*, Vol. 485, 1984.
- [Ki] D.G. Kirkpatrick, “Optimal Search in Planar Subdivisions”, *SIAM J. Comput.*, 12 (1983), pp. 28-35.
- [Ku] D.T. Kuan, “Terrain Map Knowledge Representation for Spatial Planning”, *Proc. The First Conference of Artificial Intelligence Applications*, IEEE Computer Society, Denver, CO, December 1984, pp. 578-584.
- [KBZD] D.T. Kuan, R.A. Brooks, J.C. Zamiska, and M. Das, “Automatic Path Planning for a Mobile Robot Using a Mixed Representation of Free Space”, *Proc. The First Conference of Artificial Intelligence Applications*, IEEE Computer Society, Denver, CO, December 1984, pp. 70-74.
- [KZB] D.T. Kuan, J.C. Zamiska, and R.A. Brooks, “Natural Decomposition of Free Space for Path Planning”, *Proc. IEEE International Conference on Robotics and Automation*, IEEE Computer Society, St. Louis, MO, March 1985, pp. 168-173.
- [Le] D.T. Lee, “Proximity and Reachability in the Plane”, Ph.D. Thesis, Technical Report ACT-12, Coordinated Science Laboratory, University of Illinois, Nov. 1978.
- [LC] D.T. Lee and Y.T. Ching, “The Power of Geometric Duality Revisited”, *Information Processing Letters*, **21** (1985), pp. 117-122.
- [LP] D.T. Lee and F.P. Preparata, “Euclidean Shortest Paths in the Presence of Rectilinear Boundaries”, *Networks*, 14 (1984), pp. 393-410.
- [LW] T. Lozano-Perez and M.A. Wesley, “An Algorithm for Planning Collision-Free Paths Among Polyhedral Obstacles”, *Communications of the ACM*, Vol. 22, No. 10 (1979), pp. 560-570.
- [LS] V.J. Lumelsky and A.A. Stepanov, “Path-Planning Strategies for a Point Mobile Automaton Moving Amidst Unknown Obstacles of Arbitrary Shape”, *Algorithmica*, **2**, (1987), pp. 403-430.
- [Ly] L.A. Lyusternik, *Shortest Paths: Variational Problems*, Macmillan Company, New York, 1964.
- [Meh] K. Mehlhorn, *Multi-dimensional Searching and Computational Geometry*, EATCS Monographs on Theoretical Computer Science, Springer-Verlag, Berlin, Heidelberg, New York, Tokyo, 1984.
- [Men1] A. Meng, “Free-Space Modeling and Geometric Motion Planning Under Unexpected Obstacles”, Technical Report, Texas Instruments Artificial Intelligence Lab, 1987.

- [Men2] A. Meng, "Free-Space Modeling and Path Planning Under Uncertainty for Autonomous Air Robots", Technical Report, Texas Instruments Artificial Intelligence Lab, 1987.
- [MT] M.B. Metea and J.J.-P. Tsai, "Route Planning for Intelligent Autonomous Land Vehicles Using Hierarchical Terrain Representation", *Proc. IEEE International Conference on Robotics and Automation*, Raleigh, NC, March/April 1987.
- [Mi1] J.S.B. Mitchell, "An Autonomous Vehicle Navigation Algorithm", *Proc. SPIE Applications of Artificial Intelligence*, Vol. 485, pp. 153-158, 1984.
- [Mi2] J.S.B. Mitchell, "Planning Shortest Paths", PhD Thesis, Department of Operations Research, Stanford University, August, 1986. (Available as Research Report 561, Artificial Intelligence Series, No. 1, Hughes Research Laboratories, Malibu, CA.)
- [Mi3] J.S.B. Mitchell, "Shortest Rectilinear Paths Among Obstacles", Technical Report No. 739, School of Operations Research and Industrial Engineering, Cornell University, April, 1987.
- [Mi4] J.S.B. Mitchell, "Shortest Paths Among Obstacles, Zero-Cost Regions, and Roads", Technical Report, Department of Operations Research, Cornell University, 1987. (Presented at the TIMS/ORSA meeting, May, 1987.)
- [Mi5] J.S.B. Mitchell, "On Maximum Flows In Polyhedral Domains", Extended Abstract, Department of Operations Research, Cornell University, 1987.
- [Mi6] J.S.B. Mitchell, "On the Maximum Concealment Problem", Technical Report in preparation, Department of Operations Research, Cornell University, 1988.
- [MMP] J.S.B. Mitchell, D.M. Mount, and C.H. Papadimitriou, "The Discrete Geodesic Problem", *SIAM Journal on Computing*, **16**, No. 4, pp. 647-668, August, 1987.
- [MP1] J.S.B. Mitchell and C.H. Papadimitriou, "Planning Shortest Paths", *SIAM Conference on Geometric Modeling and Robotics*, Albany, NY, July 15-19, 1985.
- [MP2] J.S.B. Mitchell and C.H. Papadimitriou, "The Weighted Region Problem", Technical Report, Department of Operations Research, Stanford University, 1986. Extended abstract appears in *Proc. Third Annual ACM Conference on Computational Geometry*, Waterloo, Ontario, June 1987. (Submitted to *Journal of the ACM*.)
- [MPK] J.S.B. Mitchell, D.W. Payton, and D.M. Keirse, "Planning and Reasoning for Autonomous Vehicle Control", *International Journal of Intelligent Systems*, **II**, (1987), pp. 129-198.
- [Mo] H.P. Moravec, "Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover", Carnegie-Mellon Robotics Institute, Pittsburgh, PA, Technical Report CMU-RI-TR-3, September, 1980.
- [Ni] N.J. Nilsson, *Principles of Artificial Intelligence*, Tioga Publishing Co., Palo Alto, CA, 1980.
- [OIRK] B.J. Oommen, S.S. Iyenger, S.V. Rao, and R.L. Kashyap, "Robot Navigation in Unknown Terrains Using Learned Visibility Graphs. Part I: The Disjoint Convex Obstacle Case", Technical Report SCS-TR-86, School of Computer Science, Carleton University, February 1986.

- [O'R] J. O'Rourke, "Finding a Shortest Ladder Path: A Special Case", Technical Report, Institute for Mathematics and its Applications, Preprint Series No. 353, University of Minnesota, October 1987.
- [OSB] J. O'Rourke, S. Suri, and H. Booth, "Shortest Paths on Polyhedral Surfaces", Manuscript, Johns Hopkins University, 1984.
- [Pa] C.H. Papadimitriou "An Algorithm for Shortest-Path Motion in Three Dimensions", *Information Processing Letters*, 20 (1985), pp. 259-263.
- [PaS] C.H. Papadimitriou and E.B. Silverberg, "Optimal Piecewise Linear Motion of an Object Among Obstacles", *Algorithmica*, 2, 1987, pp. 523-539.
- [PrS] F.P. Preparata and M.I. Shamos, *Computational Geometry*, Springer-Verlag, New York, Berlin, Heidelberg, Tokyo, 1985.
- [QFP] F.K.H. Quek, R.F. Franklin, and F. Pont, "A Decision System for Autonomous Robot Navigation Over Rough Terrain", *Proc. SPIE Applications of Artificial Intelligence*, Boston, 1985.
- [RIJW1] N.S.V. Rao, S.S. Iyengar, C.C. Jorgensen, and C.R. Weisbin, "Robot Navigation in an Unexplored Terrain", *Journal of Robotic Systems*, 3, No. 4, (1986), pp. 389-407.
- [RIJW2] N.S.V. Rao, S.S. Iyengar, C.C. Jorgensen, and C.R. Weisbin, "On Terrain Acquisition by a Finite-Sized Mobile Robot in Plane", *Proc. IEEE International Conference on Robotics and Automation*, Raleigh, NC, March/April 1987, pp. 1314-1319.
- [RS] J.H. Reif and J.A. Storer, "Shortest Paths in Euclidean Space with Polyhedral Obstacles", Technical Report CS-85-121, Computer Science Department, Brandeis University, April, 1985.
- [Ri] R.F. Richbourg, "Solving a Class of Spatial Reasoning Problems: Minimal-Cost Path Planning in the Cartesian Plane", Ph.D. Thesis, Naval Postgraduate School, Monterey, CA, June, 1987.
- [RRZ] R.F. Richbourg, N.C. Rowe, and M.J. Zyda, "Exploiting Capability Constraints To Solve Global, Two-Dimensional Path Planning Problems", Technical Report NPS-86-006, Department of Computer Science, Naval Postgraduate School, 1986.
- [RRZM] R.F. Richbourg, N.C. Rowe, M.J. Zyda, and R. McGhee, "Solving Global Two-Dimensional Routing Problems Using Snell's Law and A* Search", *Proc. IEEE International Conference on Robotics and Automation*, Raleigh, NC, March/April 1987, pp. 1631-1636.
- [Ro] N.C. Rowe, "Roads, Rivers, and Rocks: Optimal Two-Dimensional Route Planning Around Linear Features for a Mobile Robot", Technical Report NPS52-87-027, Computer Science Department, Naval Postgraduate School, June 1987.
- [RR] N.C. Rowe and R.F. Richbourg, "A New Method for Optimal Path Planning Through Nonhomogeneous Free Space", Technical Report NPS52-87-003, Computer Science Department, Naval Postgraduate School, June 1987.
- [RW] K.D. Rueb and A.K.C. Wong, "Structuring Free Space as a Hypergraph for Roving Robot Path Planning and Navigation", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9, No. 2, March 1987, pp. 263-273.

- [SW] H. Samet and R.E. Webber, "Storing a Collection of Polygons Using Quadtrees", *ACM Transactions on Graphics*, 4, No. 3, July 1985, pp. 182-222.
- [SSH] J.T. Schwartz, M. Sharir, and J. Hopcroft, Ed., *Planning, Geometry, and Complexity of Robot Motion*, Ablex Series in Artificial Intelligence, Ablex Publishing Corporation, Norwood, NJ, 1987.
- [SY] J.T. Schwartz and C. Yap, Ed., *Algorithmic and Geometric Aspects of Robotics, Vol. 1*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1987.
- [SS] M. Sharir and A. Schorr, "On Shortest Paths in Polyhedral Spaces", *SIAM Journal of Computing* Vol. 15, No. 1, pp. 193-215, February 1986.
- [Sm] T. Smith, Private communication, Department of Computer Science, University of California, Santa Barbara, 1986.
- [So] D. Souvaine, "Computational Geometry in a Curved World", Ph.D. Thesis, Technical Report CS-TR-094-87, Department of Computer Science, Princeton University, October, 1986.
- [St] H.I. Stern, "A Routing Algorithm for a Mobile Search Robot", in *Impacts of Microcomputers on Operations Research*, Gass, Greenburg, Hoffman, and Langley, Eds., Elsevier Science Publishing Co., 1986, pp. 73-87.
- [Su] I.E. Sutherland, "A Method for Solving Arbitrary Wall Mazes by Computer", *IEEE Transactions on Computers*, C-18, No. 12, (1969), pp. 1092-1097.
- [TV] R.E. Tarjan and J. Van Wyk, "An $O(n \log \log n)$ Algorithm for Triangulating Simple Polygons", Technical Report, Bell Labs, 1986.
- [Th1] C.E. Thorpe, "Path Relaxation: Path Planning for a Mobile Robot", *Proc. AAAI National Conference on Artificial Intelligence*, Austin, TX, August 6-10, 1984, pp. 318-321.
- [Th2] C.E. Thorpe, "FIDO: Vision and Navigation for a Robot Rover", Ph.D. Thesis, Technical Report CMU-CS-84-168, Department of Computer Science, Carnegie-Mellon University, 1984.
- [We] E. Welzl, "Constructing the Visibility Graph for n Line Segments in $O(n^2)$ Time", *Information Processing Letters*, Vol. 20 (1985), pp. 167-171.
- [WWW] P. Widmayer, Y. F. Wu, and C. K. Wong, "Distance Problems in Computational Geometry with Fixed Orientations", *Proc. First Annual ACM Conference on Computational Geometry*, 1985.
- [Zi] K. Zikan, Private communication, Department of Operations Research, Stanford University, Stanford, CA, 1986.