

# Path Planning In Image Space For Autonomous Robot Navigation in Unstructured Environments

---

**Michael W. Otte**

Department of Computer Science  
University of Colorado at Boulder  
Boulder, CO 80309-0430  
michael.otte@colorado.edu

**Scott G. Richardson**

Department of Computer Science  
University of Colorado at Boulder  
Boulder, CO 80309-0430  
scottric@gmail.com

**Jane Mulligan**

Department of Computer Science  
University of Colorado at Boulder  
Boulder, CO 80309-0430  
jane.mulligan@colorado.edu

**Gregory Grudic**

Department of Computer Science  
University of Colorado at Boulder  
Boulder, CO 80309-0430  
gregory.grudic@colorado.edu

## Abstract

Path planning systems using graph-search algorithms such as A\* usually operate in uniform plan-view occupancy grids. However, the sensors used to construct these grids observe the environment in their own sample space based on sensor type and viewpoint. In this paper we present an *image space* technique for path planning in unknown unstructured outdoor environments. Our method differs from previous techniques in that we perform path search directly in image space—the native sensor space of the imaging sensor. After an image space path has been found, it is used for navigation in the real world. By operating at the resolution of the image sensor, image space planning facilitates accurate robot vs. obstacle localization, and enables a high degree of

movement precision. Our image space planning techniques can potentially be used with many different kinds of sensor data, and we experimentally evaluate the use of stereo disparity and color information. We present an extension to the basic image space planning system called the *cylindrical planner* that simulates a  $2\pi$  field-of-view with a cylindrically shaped occupancy grid. We believe image space planning is well suited for use in the local subsystem of a hierarchical planner, and implement a *hybrid hierarchical planner* that utilizes the cylindrical planner as a local planning subsystem and a 2-dimensional Cartesian planner as the global planning subsystem. All three systems are implemented and experimentally tested on a real robot. We evaluate the failure modes of image space planning, and discuss how to avoid them. We find that image space enables precise real-time near-field planning.

## 1 Introduction

Autonomous robot navigation is the search for a series of actions that move a robot from an initial state to a goal state. Typically, an internal representation of the world is used to select actions from a movement set defined by the robot's kinematics. For autonomous navigation in unstructured environments, the world representation is often encoded in a map. Navigation is achieved by finding a map-based path between map-based start and goal locations, extrapolating that path to the real world, and then following it. Occupancy grid maps are discretized environmental representations in which each grid cell stores data about a particular piece of the real world [Herman, 1986, Elfes, 1989, Goldberg et al., 2002, Kolski et al., 2006]. Occupancy grids are traditionally uniform 2D or 3D Cartesian models containing traversability data (e.g. cost or obstacle probability). A path is found between the occupancy grid start and goal locations using a graph-search algorithm [Dijkstra, 1959, Hart et al., 1968, Stentz, 1995, Krishnaswamy and Stentz, 1995, Koenig and Likhachev, 2002, Ferguson and Stentz, 2006], a finite element potential field model [Khatib, 1986, Koren and

Borenstein, 1991], or a combination of the two [Murray and Little, 1998]. We focus primarily on the former paradigm, where variants of A\* and D\* are commonly used.

In unknown environments, the occupancy grid is not known *a priori* and must be created as the robot observes the world with its sensors. Computer vision methods like 3D stereo reconstruction can be used to obtain depth information [Murray and Jennings, 1998, van den Mark et al., 2001, Matthies, 1992]. Map building techniques that project stereo or laser depth data into a Cartesian model of the world depend on pose estimates that may be inaccurate in the short term. We investigate using the sensor image to guide robot motion without projection, particularly for real-time near-field obstacle avoidance.

Planning in the space of image sensor samples, or *image space*, presents several challenges. Due to a sensor's pose and projection, the sensor sample array does not represent a uniform grid in the world; i.e., near- and far-field samples arise from regions of differing size (several square centimeters to meters, respectively). We exploit occupancy grid planning methods in image space by using nonuniform graph-search edge costs. Once an image space path is computed, motion control is defined based on the first-person relative position of a near-field path waypoint. Since the image changes with the motion of the robot, the use of memory offers interesting trade-offs. In order to avoid oscillation as obstacles move in and out of the field-of-view, we present a method for accumulating image data in a cylindrical panorama as the robot moves. We call this type of system a *cylindrical planner*.

A *hierarchical planner* (as considered here) is a planning system that divides the path planning problem into the two parallel tasks of local planning and global planning. The global planner finds an approximate path to the goal based on a coarse world representation, while the local planner is concerned with obstacle avoidance and navigation toward sub-goals (provided by the global planner) through a more precise world representation. A hierarchical planner allows the global map to expand as the robot encounters new parts of the environment, while facilitating time-bounded obstacle avoidance in the near field via the local

map. A diverse variety of hierarchical planning systems have been proposed in the literature [Sugiyama et al., 1994, Chen, 1990, Krogh and Thorpe, 1986, Hong et al., 1991, Gat et al., 1990, Carsten et al., 2007] and many others. A common implementation uses two 2D Cartesian occupancy grids. The local planner maintains a fixed-size high-resolution map that is centered on the robot, while the global planner accumulates a low-resolution map that expands with exploration and is anchored to a global frame of reference. We propose a *hybrid hierarchical planner* that performs local navigation in a cylindrical planner and global navigation in a 2D Cartesian occupancy grid. We believe the strengths of image planning (discussed in Section 2.1) justify its use as a local planner, while Cartesian space is better suited to long-distance planning tasks requiring persistent memory. A hybrid hierarchical planner combines the strengths of both planning paradigms in a single system.

This paper is organized as follows: Section 2 contains an overview of image planing, a survey of related work, and a discussion on cost maps motivating our method. The basic image planner is introduced in Section 3, and the cylindrical and hybrid hierarchical planners are described in Sections 4 and 5, respectively. In Section 6 we outline our hardware, and in Section 7 we perform a series of experiments to evaluate the image planner, the cylindrical planner, and the hybrid hierarchical planner against a state-of-the-art 2D Cartesian hierarchical planer. Specific system parameter values are also stated in Section 7. Results are presented in Section 8 and conclusions in Section 9.

## 2 Background

### 2.1 Image space path planning

A *grid-based sensor* or *image sensor* captures multiple simultaneous observations about the world via a grid of sensor elements [Sonka et al., 2008]. An *image* is the instantaneous output of an image sensor and can be represented as a matrix, the elements of which are called *pixels*. *Image space* is a coordinate space of the sensor grid with unit vectors defined by pixel

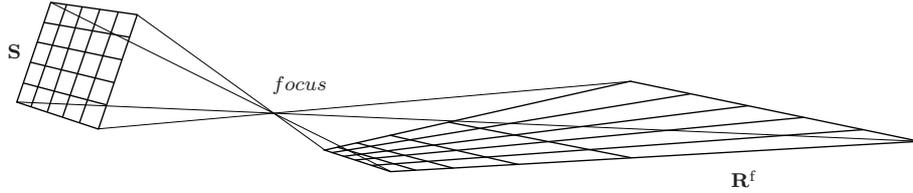


Figure 1: A projection of pixels from image space  $\mathbf{S}$  onto a flat level groundplane  $\mathbf{R}^f$ .

height and width. In *image space planning* a path through an image is used for navigation. Digital cameras are commonly used for image space planning. We experiment with stereo disparity and RGB (red, green, blue) intensity values—both captured with cameras—and refer to image sensors as cameras. However, our general methods can also be used with other grid-based sensors (e.g. LIDAR, SONAR, infrared, etc).

We restrict our discussion to scenarios where a passive image sensor is mounted on an autonomous robot. We assume that spatial transformations are known between the sensor and robot coordinate frames, that the sensor-grid observes the ground-surface (or obscuring obstacles), and that the groundplane parameters  $(\mathbf{n}, d)$  are known relative to the calibrated camera frame. Only one real-world path can be described by projecting an image space path onto the real-world ground-surface. We assume that connectivity along the image space path implies connectivity along that real-world path. Finally, we assume basic prior knowledge about traversability given a sensor type (e.g. stereo disparity or color), and also a known goal location on the ground-surface.

Image space preserves the accuracy and precision of sampled scene information given an instantaneous sensor reading. Consequently, image space planning minimizes the number of map locations needed to accurately represent a scene projection captured by an image sensor. Map size is fixed by the properties of the sensor, which guarantees real-time obstacle avoidance by placing upper-bounds on the time required for graph-search. Due to foreshortening and perspective projection, near-field pixels capture scene information at a higher resolution than far-field pixels, although both are represented at the finest granularity available (Figure 1). As a result, image space provides robust localization of near-field

obstacles, enabling navigation through tight passages and cluttered environments.

If other information is available, in addition to the current image, then image space may not provide the most accurate world representation. For example, there is a point beyond which a grid-based Cartesian map will have a higher resolution than an image space map. This is because the granularity of a Cartesian map is fixed, relative to the real world, while the granularity of an image space map becomes coarser toward the horizon (Figure 1). In general, image space projections cannot represent what is behind view obstructions. This problem is initially encountered by any system that must discover the world. However, image space lacks the ability to remember previously explored areas that are currently outside the field-of-view. We discuss this problem, including solutions, in Section 3.4.4.

## 2.2 Related work

The most closely related work to our own has been pursued by [Ollis et al., 2008]. In this work, image space and Cartesian planners operate in parallel. If a low-cost image-based path is found, then the image planner guides the default Cartesian planner by providing a waypoint from the image-based path. This technique is used in conjunction with cost from color data to find paths beyond the range of stereo disparity. In contrast, our approach focuses on high-resolution near-field planning.

Autonomous highway driving algorithms rely on image features such as lane markings and road color/texture data to infer knowledge about a scene [Tsugawa et al., 1979, Thorpe et al., 1988, Pomerleau, 1989, Jochem et al., 1995, Mateus et al., 2005]. A related off-road system assumes prior vehicle tire tracks or pedestrian footprints [Song et al., 2007]. However, these features are not guaranteed to exist, or be useful, in all outdoor environments.

Visual servoing methods use the appearance relationship between a camera's current field-of-view and a predefined target image to calculate steering commands [Feddema and Mitchell, 1989, Hutchinson et al., 1996, Gaussier et al., 1997, Winters et al., 2000, Cowan et al., 2002,

Matsumoto et al., 2000, Zhang and Ostrowski, 2002, Vidal et al., 2003]. Navigation through unknown terrain implies we cannot obtain a predefined image of the goal state.

Reactive local planners often use polar models to determine navigational headings. [Borenstein and Koren, 1991] create a polar histogram from data accumulated in a global Cartesian map, and [Minguez and Montano, 2004] build a polar map from a ring of depth sensors. In either case, the local map is a 1D vector and each element represents the cost of moving in a particular direction. These techniques are similar to our cylindrical planner in that they utilize polar coordinates. However, the additional dimension of the cylindrical planner enables it to formulate an entire path between the robot and goal.

Range data from stereo vision is frequently used in conjunction with Cartesian planning systems [Murray and Little, 1998, van den Mark et al., 2001, Goldberg et al., 2002, Sabe et al., 2004, Carsten et al., 2007]. These approaches project data into top-down or 3D Cartesian representations of the world, instead of planning directly in image space.

### 2.3 Cost vs. distance

Each cell or *map element* in an occupancy grid map represents a unique subset of the real world, or *world element*, and has a value associated with the cost of movement through the corresponding world element. Map elements are usually arranged in rows and columns of unit-length squares. This configuration can be stored as a 2D array  $\mathbf{G}$ , where  $\mathbf{G}_{n,m}$  maintains the cost associated with the row  $n$  column  $m$  map element. If a graph structure is imposed on the map, then a path between two map locations can be found with a graph-search algorithm. The former is typically achieved by recasting map elements as graph nodes, and then creating edges from each node to its neighbors. Common structures include the 4- and 8-connected graphs. Let the graph formalization of the cost map be denoted  $\mathbf{G}^\Gamma$  to differentiate it from the map array  $\mathbf{G}$ . The graph node  $\mathbf{G}_{n,m}^\Gamma$  is associated with the row  $n$  and column  $m$  map element. Graph search algorithms operate on edges, not nodes; therefore,

a method of determining edge cost must be devised. This is often done by defining edge cost to be the map cost value  $\mathbf{G}_{n,m}$  or  $\mathbf{G}_{k,j}$  associated with the nodes  $\mathbf{G}_{n,m}^\Gamma$  and  $\mathbf{G}_{k,j}^\Gamma$  that an edge connects, or the average value of  $\mathbf{G}_{n,m}$  and  $\mathbf{G}_{k,j}$ .

Consider a uniform world such that all elements in a 2D top-down Cartesian map have identical cost. Let the map graph be 8-connected. The Euclidean distance between neighboring world element centers is given, to a scale factor, by the distance between the corresponding map elements. The scale factor can be ignored without affecting search results, so we assume it is 1. If edge cost is defined as the map cost  $\mathbf{G}_{n,m}$  of the node  $\mathbf{G}_{n,m}^\Gamma$  that the edge is directed to (or from or their mean), then it costs 1 to move between neighboring nodes. An edge cost of 1 reflects the 1 Euclidean distance between horizontal and vertical neighbors, but is inaccurate with respect to the  $\sqrt{2}$  distance between diagonal neighbors. Edge cost should reflect both map cost and real world distance. Therefore, we redefine edge cost to be *map cost integrated over the Euclidean distance between map elements*. In a discretized map this reduces to map cost multiplied by the distance through a map element.

The idea of integrating cost over distance is used in [Ferguson and Stentz, 2006] and elsewhere; however, we have found little discussion in the literature as to why this is a good idea. In fact, it is supported by a geometric argument and a physical analogy. (Geometric) Let point cost be a measure of traversability associated with a point. Let the integration of point cost over a grid element's area yield that grid element's map cost. Grid elements are assumed to be unit-length squares of uniform point cost; thus, point cost is equal to the map cost of the grid element containing a particular point. Edges are defined as line segments spanning neighboring grid element centers, and one-dimensional line cost is found by integrating point cost over a segment's length. Thus, edge cost is map cost integrated over Euclidean distance. (Physical) Assume map cost is a force opposing movement within a map element. The most efficient way through a map is along the path of minimum work. Work is force integrated over distance; therefore, the most efficient path is found by minimizing map cost integrated over Euclidean distance.

To disambiguate between map cost (i.e. point cost, force) and edge cost (i.e. line cost, work), we employ the terms used in the physical analogy—we refer to map cost as force and edge cost as work. The concepts of force, distance, and work are integral to our image space planning system. Due to perspective, image space world elements are not square and vary in size from the near- to far-field. We hypothesize that image space paths will be more useful for navigation if the distance quantity reflects real world spatial relationships.

## 2.4 Memory and myopic behavior

Let *rotational memory* and *translational memory* denote, respectively, an ability to remember map data after an arbitrarily large rotation or translation. Local Cartesian maps do not have translational memory because data is forgotten after a translation beyond the map’s fixed size. They do have rotational memory because the robot can pivot without losing information. Global Cartesian maps have both rotational and translational memory. Lack of translational or rotational memory can induce myopic behavior in the form of *translational* or *rotational oscillation*, as the robot repeatedly moves away from, forgets, and then returns to a goal blocking obstacle. Myopic behavior is not the only cause of oscillation; e.g. [Carsten et al., 2007] observe oscillation from conflicting votes in a vote-based planner.

## 3 The basic image planner

We call our basic image space planner the *image planner*. It creates an image space occupancy grid from pixel data, and makes no additional assumptions to those in Section 2.1.

### 3.1 Occupancy grids and cost functions in image space

Let  $\mathbf{S}$  be the  $h$  row and  $w$  column output of an image sensor. Pixels are denoted  $\mathbf{S}_{n,m}$ , where  $1 \leq n \leq h$  and  $1 \leq m \leq w$ , and  $n = 1$  and  $m = 1$  correspond to the top row and left column, respectively.  $\mathbf{O}$  is the image space occupancy grid built from  $\mathbf{S}$ . Map values are determined

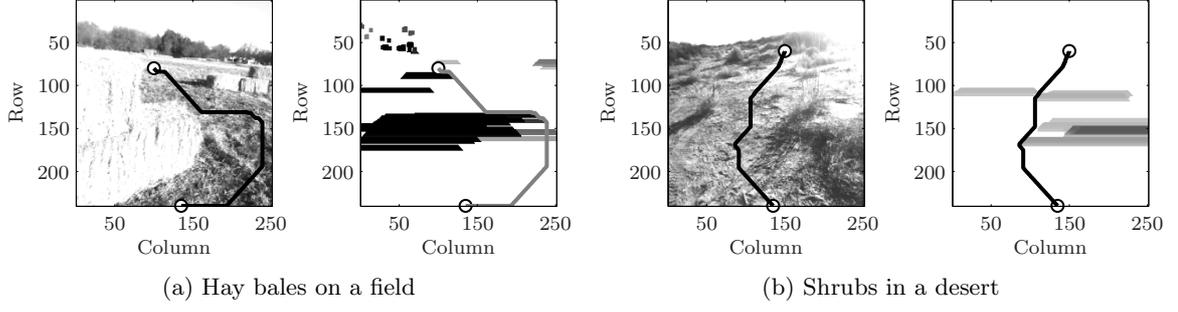


Figure 2: Image space paths through grayscale images and corresponding occupancy grids.

as a function of the corresponding image pixel data, and in the most basic implementation there is a one-to-one and onto mapping from pixels to grid locations.

$$\mathbf{O}_{n,m} = f(\mathbf{S}_{n,m})$$

Let  $\mathbf{O}^\Gamma$  be the 8-connected graph of  $\mathbf{O}$ . Node  $\mathbf{O}_{n,m}^\Gamma$  is associated with map element  $\mathbf{O}_{n,m}$ , and is a neighbor of  $\mathbf{O}_{k,j}^\Gamma$  for all  $k = n \pm 1$  and  $j = m \pm 1$ , where  $1 \leq k \leq h$  and  $1 \leq j \leq w$ . Let arc  $i$  connect  $\mathbf{O}_{k,j}^\Gamma$  to  $\mathbf{O}_{n,m}^\Gamma$ . The length of  $i$  is  $D_i$ . We interpret  $\mathbf{O}_{n,m}$  as a force  $F_i$  impeding progress along  $i$ , and define the edge cost of  $i$  as work  $W_i$ .

$$W_i = F_i D_i$$

Let  $\mathbf{R}$  denote the real world. To move from the current world location  $\mathbf{R}_s$  to a goal location  $\mathbf{R}_g$ , the system transforms  $\mathbf{R}_s$  and  $\mathbf{R}_g$  into their image space graph equivalents  $\mathbf{O}_s^\Gamma$  and  $\mathbf{O}_g^\Gamma$ , respectively, and then uses a variant of the A\* algorithm to find the minimum work path  $P_{\mathbf{O}}^{\min}$  between  $\mathbf{O}_s^\Gamma$  and  $\mathbf{O}_g^\Gamma$ .

$$W_{P_{\mathbf{O}}^{\min}} = \min_{P_{\mathbf{O}}} \left( \sum_{i \in P_{\mathbf{O}}} F_i D_i \right)$$

The optimality of A\* is only insured if  $f(\mathbf{S}_{n,m})$  returns nonnegative values over the entire range of  $\mathbf{S}_{n,m}$  allowed by the sensor; thus,  $f(\mathbf{S}_{n,m})$  is system dependent. Figure 2 shows two examples of  $P_{\mathbf{O}}^{\min}$  projected into  $\mathbf{O}$  and the corresponding grayscale images.

$\mathbf{O}_s^\Gamma$  represents the robot's projected location in the image  $\mathbf{S}$ , and depends on the relationship between the robot and the camera. If the camera is forward facing and fixed in the center of the robot, then  $\mathbf{O}_s^\Gamma$  is the bottom center graph node  $\mathbf{O}_{h, \lfloor w/2 \rfloor}^\Gamma$ . The method used to position

$\mathbf{O}_g^\Gamma$  is task dependent. A system designed to seek a specific color/pattern target may extract the coordinates of  $\mathbf{O}_g^\Gamma$  directly from an RGB image. Goals defined by GPS coordinates must be projected into image space, implying the existence of a transformation function. The robot is programmed to rotate toward  $\mathbf{R}_g$  whenever  $\mathbf{R}_g$  is outside the field-of-view (FOV).

### 3.2 Force metrics

The map value  $\mathbf{O}_{n,m}$  is a function of pixel data  $\mathbf{S}_{n,m}$  and determines the force  $F_i$  required to move along arc  $i$  toward node  $\mathbf{O}_{n,m}^\Gamma$ .

$$F_i = \mathbf{O}_{n,m} = f(\mathbf{S}_{n,m}) \quad (1)$$

Equation 1 is a force metric defining the relationship between sensor data and force in  $\mathbf{O}^\Gamma$ . To be useful,  $f(\mathbf{S}_{n,m})$  should be representative of real world traversability. Obstacles should have high (or lethal) values relative to flat terrain. We experiment with two force metrics. The first metric  $F_i^d$  assumes that pixel  $\mathbf{S}_{n,m}$  contains disparity information. The second metric  $F_i^c$  assumes that  $\mathbf{S}_{n,m}$  is a triple of RGB intensity values.

#### 3.2.1 Force from stereo disparity

Given two images of a scene from different viewpoints, stereo disparity is the difference between the image locations to which a point is projected. Given rectified images and cameras of equal focal length, disparity  $\mathbf{S}_{n,m}$  is related to depth  $z(\mathbf{S}_{n,m})$  by:

$$z(\mathbf{S}_{n,m}) = \frac{c_f c_b}{\mathbf{S}_{n,m}}$$

where  $c_f$  and  $c_b$  are focal length and the baseline between optical centers, respectively [Sonka et al., 2008]. Disparity approaches zero as depth approaches infinity and vice-versa. However, given a finite camera resolution, disparity may go to zero at less than 30 meters for practical  $w$ ,  $h$ ,  $c_b$ , and  $c_f$ . We define a force metric from disparity as follows:

$$F_i^d = \mathbf{O}_{n,m} = 1 + c_{scl} \left| \mathbf{S}_{n,m} - \mathbf{S}_{n,m}^f \right| \quad (2)$$

where  $\mathbf{S}_{n,m}^f$  is the disparity observed from a known flat level groundplane  $\mathbf{R}^f$ , and  $c_{\text{scl}}$  is a scaling constant.  $c_{\text{scl}}$  is tuned so that  $F_i^d$  will be greater than a threshold  $c_{\text{thd}}$  for obstacles within stereo range. A concept of lethality is added by explicitly resetting  $F_i^d$  to  $10^8 c_{\text{thd}}$  if  $F_i^d$  is greater than  $1.5c_{\text{thd}}$ . Distance may be used as an admissible heuristic for estimating work during graph-search because  $F_i^d \geq 1$ .

### 3.2.2 Force from color

Our color data is composed of three dimensions, red  $\mathbf{S}_{n,m}^{\text{red}}$ , green  $\mathbf{S}_{n,m}^{\text{grn}}$ , and blue  $\mathbf{S}_{n,m}^{\text{blu}}$ . Values range between 0 and 255. In practice, a force from color metric can be provided by a machine-learning sub-system that is trained on-line [Grudic et al., 2007, Procopio et al., 2007, Ollis et al., 2008]. We implement a predefined metric  $F_i^c$  as a proof of concept. Any predefined force from color metric is only valid in a subset of environments, and the simple technique described here is even brittle in a single environment. Hand labeled training examples, one each for traversable terrain and obstacles, are provided to the system prior to a run. The color dimension of largest magnitude is denoted *gnd* or *obs* for the traversable or obstacle example, respectively,  $\text{gnd}, \text{obs} \in \{\text{red}, \text{grn}, \text{blu}\}$ . This method of determining an example’s associated color dimension is simple and works in practice, but it may not be superior to any other.  $\hat{\mathbf{S}}$  is a normalized version of  $\mathbf{S}$  with values between 0 and 1.  $\hat{\mathbf{S}}_{n,m}^{\text{obs}}$  and  $\hat{\mathbf{S}}_{n,m}^{\text{gnd}}$  are normalized intensity values along the *obs* and *gnd* dimensions, respectively.

$$F_i^c = \mathbf{O}_{n,m} = 1 + c_{\text{scl}} \left( \hat{\mathbf{S}}_{n,m}^{\text{obs}} - \hat{\mathbf{S}}_{n,m}^{\text{gnd}} \right) \quad (3)$$

$F_i^c$  is scaled between 1 and  $c_{\text{thd}}$  using  $c_{\text{scl}}$ , and defined as lethal if greater than  $c_{\text{thd}}/3$ . Distance may be used as an admissible heuristic for estimating work during graph-search because  $F_i^c \geq 1$ .  $F_i^c$  fails if  $\text{obs} = \text{gnd}$ , and is only valid in environments where obstacles and traversable terrain differ in color. It is worth reiterating that  $F_i^c$  is used as a proof of concept. Our goal is to show that the image planner can navigate using color data.

### 3.3 Distance metrics

$D_i$  is defined to be the length of arc  $i$  from node  $\mathbf{O}_{k,j}^\Gamma$  to node  $\mathbf{O}_{n,m}^\Gamma$  in the graph  $\mathbf{O}^\Gamma$ . In general,  $D_i$  is a function of grid coordinates.

$$D_i = f(k, j, n, m)$$

$D_i$  must be nonnegative for the optimality of  $A^*$ , and nonzero to insure that  $A^*$  is useful ( $A^*$  ignores all force information and degenerates into an uninformed breadth-first search when  $D_i = 0$ ). We now discuss two distance metrics used in our experiments.

#### 3.3.1 The $L^2$ norm in image space

In 2D top-down Cartesian systems, the  $L^2$  norm in map space is a constant multiple of the  $L^2$  norm in the real world (assuming elevation can be ignored), and it makes sense to define arc length as the Euclidean distance between grid element centers. We experiment using the  $L^2$  norm of image space as a distance metric  $D_i^{\mathbf{S}}$ .

$$D_i^{\mathbf{S}} = \sqrt{(n - k)^2 + (m - j)^2} \quad (4)$$

$D_i^{\mathbf{S}}$  is the distance between  $\mathbf{O}_{k,j}$  and  $\mathbf{O}_{n,m}$ , assuming that grid element side lengths are 1. The distance between horizontal, vertical, and diagonal neighbors is 1, 1, and  $\sqrt{2}$ , respectively.  $D_i^{\mathbf{S}}$  does *not* accurately model the distances between image space world elements—which increase toward the horizon.

#### 3.3.2 The flat world distance metric

We define the flat world distance metric  $D_i^{\mathbf{R}^f}$  to be the Euclidean distance between the two points  $\mathbf{x}_1$  and  $\mathbf{x}_2$  on a flat groundplane  $\mathbf{R}^f$  that are projected to the centers of pixels  $\mathbf{S}_{k,j}$  and  $\mathbf{S}_{n,m}$ , respectively (Figure 3). To calculate the positions of  $\mathbf{x}_1$  and  $\mathbf{x}_2$ , we find the intersection of  $\mathbf{R}^f$  and the rays extending from pixel centers through the camera’s optical center. These are nominal world grid distances, assuming traversable regions of the scene lie

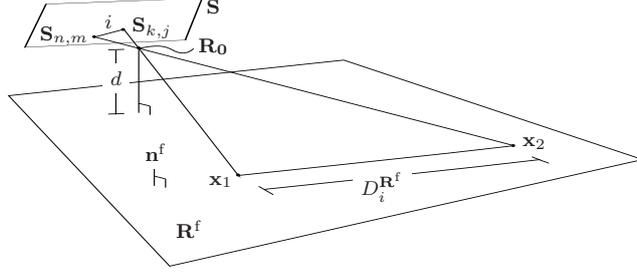


Figure 3: A 3D view of the relationship between image space arc  $i$  and the distance  $D_i^{\mathbf{R}^f}$ .

on  $\mathbf{R}^f$ . Tracing pixel rays to the groundplane reverses the projection process by following the line of sight from the pixel. Given the known camera intrinsic matrix  $\mathbf{K}$  we convert from pixel coordinates to normalized coordinates  $\mathbf{R}_d$ .

$$\mathbf{R}_d = \mathbf{K}^{-1} [m, n, 1]^T$$

The rays emanate from the camera center  $\mathbf{R}_0$  which is the origin of the camera coordinate frame. We assume the groundplane parameters  $\mathbf{R}^f = (\mathbf{n}^f, d)$  are known in the camera coordinate frame.  $\mathbf{n}^f = [n_x^f, n_y^f, n_z^f]^T$  is the plane normal and  $d$  is the perpendicular distance from the origin to the plane.  $n_x^f x + n_y^f y + n_z^f z = d$  for points on the plane described by the vector  $\mathbf{x} = [x, y, z]^T$ . To find the intersection such that  $\mathbf{R}_0 + t\mathbf{R}_d$  lies on the plane, we need to solve  $\mathbf{n}^f \cdot (\mathbf{R}_0 + t\mathbf{R}_d) - d = 0$  for  $t$ . Recall that  $\mathbf{R}_0 = [0, 0, 0]^T$  is the origin.

$$t = \frac{\mathbf{n}^f \cdot \mathbf{R}_0 + d}{\mathbf{n}^f \cdot \mathbf{R}_d} = \frac{d}{\mathbf{n}^f \cdot \mathbf{R}_d}$$

We can compute the desired intersection points using  $\mathbf{x} = t\mathbf{R}_d$ . Finally,  $D_i^{\mathbf{R}^f}$  is computed as the Euclidean distance between  $\mathbf{x}_1$  and  $\mathbf{x}_2$ .

$$D_i^{\mathbf{R}^f} = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2} \quad (5)$$

$D_i^{\mathbf{R}^f}$  defines arc lengths in  $\mathbf{O}^f$  as the flat level ground-surface distances they represent. Although the ground-surface in unstructured outdoor environments may not be flat or level, we hypothesize that  $D_i^{\mathbf{R}^f}$  more accurately models the world than  $D_i^{\mathbf{S}}$ . Movement through  $\mathbf{O}^f$  becomes easier toward the bottom of the graph because  $D_i^{\mathbf{R}^f}$  increases toward the horizon line.  $D_i^{\mathbf{R}^f}$  also reflects the fact that image space world elements become increasingly elongated toward the horizon (see Figure 1).

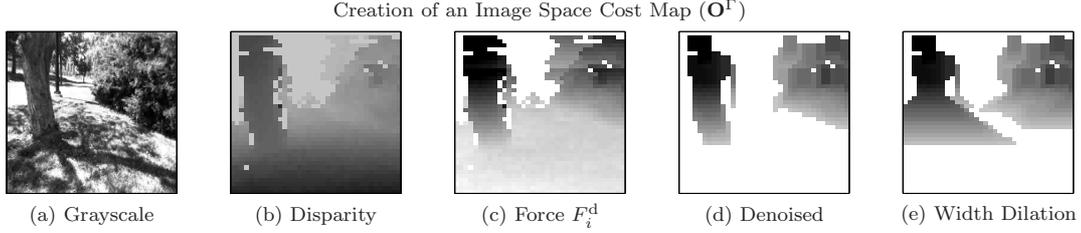


Figure 4: A grayscale scene image, the corresponding stereo disparity, and various stages of the  $F_i^d$  cost map. Light to dark corresponds to high to low disparity and cost, respectively.

### 3.4 Preprocessing

#### 3.4.1 Noise reduction

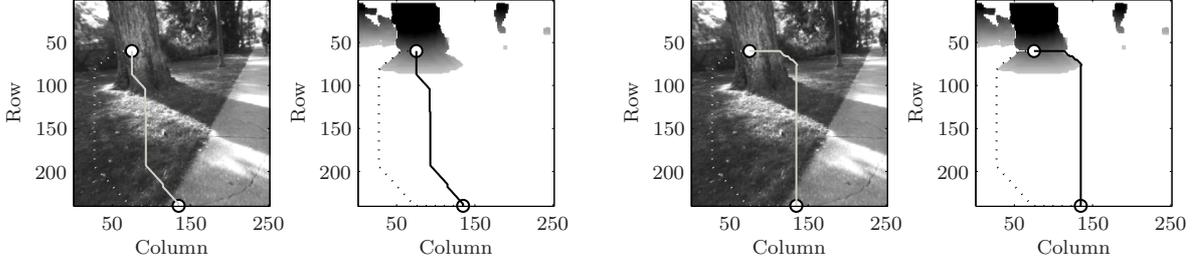
Small regions of noise are observed with both stereo and color data. We erode and then dilate  $\mathbf{O}$  to remove errors introduced by this noise. The same mask is used for both operations. To eliminate small variations in nearly-flat terrain, we reset map values less than a threshold  $c_{\text{ft}}$  to 1. Figures 4c and 4d show  $\mathbf{O}$  before and after noise reduction, respectively.

#### 3.4.2 Obstacle width dilation

The A\* search algorithm assumes the robot is a point particle. As suggested by [Kolski et al., 2006, Sugiyama et al., 1994, Kelly, 1994], we increase obstacle width in the occupancy grid by performing a dilation of radius  $\alpha$ , where  $\alpha$  is the map space equivalent of one half robot width  $c_{\text{wth}}/2$  plus a buffer  $c_{\text{buf}}$ . This ensures that paths through  $\mathbf{O}^\Gamma$  will cause the robot to avoid obstacles by at least  $c_{\text{buf}}$ . The dilation equation is given by:

$$\mathbf{O}_{n,m} = \max_j (\mathbf{O}_{n,m \pm j})$$

where  $[-\alpha] \leq j \leq [\alpha]$  and  $1 \leq (m \pm j) \leq w$ . The apparent robot width in the image  $\mathbf{S}$ , and therefore  $\mathbf{O}$ , varies due to perspective. Our dilation function assumes a flat level plane  $\mathbf{R}^f$  and approximates each pixel as an equal portion of  $\theta_w$ , where  $\theta_w$  is the horizontal FOV. Let  $D_{0 \rightarrow n,m}^{\mathbf{R}^f}$  be the distance along  $\mathbf{R}^f$  from the robot to the projected center of  $\mathbf{O}_{n,m}$ . Equation 5



(a) It requires too much work for  $D_i^{\mathbf{S}}$  to go around the tree because the planner has no notion of going behind obstacles, this results in failure.

(b) The occupancy grid has been preprocessed by setting force in the goal row equal to the minimum force value. This creates an artificial conception of obstacle depth that enables both paths to go around the tree.

Figure 5: Image space paths based on the  $D_i^{\mathbf{S}}$  and  $D_i^{\mathbf{R}^f}$  distance metrics (solid and dotted lines, respectively), through grayscale scene images and the corresponding occupancy grids.

can be used to calculate  $D_{0 \rightarrow n, m}^{\mathbf{R}^f}$  by letting  $\mathbf{x}_1 = [0, 0, -d]^T$ .  $\alpha$  is a function of  $D_{0 \rightarrow n, m}^{\mathbf{R}^f}$ .

$$\alpha = \frac{w}{\theta_w} \arcsin\left(\frac{c_{\text{wth}}/2 + c_{\text{buf}}}{D_{0 \rightarrow n, m}^{\mathbf{R}^f}}\right) \quad (6)$$

$\alpha$  can be calculated off-line. The flat level plane assumption provides a useful approximation to reality. An alternative is to perform dilation as a function of depth, assuming depth data exists (an idea not explored in this paper). Figures 4d and 4e show  $\mathbf{O}$  before and after robot width dilation, respectively.

### 3.4.3 Enabling rotation

If  $\mathbf{S}$  does not capture the area immediately adjacent to the robot, then obstacles in the bottom grid row  $\mathbf{O}_{h, 1 \dots w}$  can prevent rotation, despite the existence of maneuvering space. We explicitly allow rotation by defining the bottom row to have the minimum force value 1.

$$\mathbf{O}_{h, 1 \dots w} = 1$$

### 3.4.4 Accounting for finite obstacle depth

If the goal is a coordinate in  $\mathbf{R}$ , then a problem can arise when it is obscured by an obstacle. The goal will appear inside the obstacle in  $\mathbf{O}^{\Gamma}$ , requiring the path to penetrate the obstacle.

The cheapest path may go ‘up’ the obstacle in  $\mathbf{O}^\Gamma$ , leading the robot to collide with the obstacle in  $\mathbf{R}$  (see Figure 5a). Obstacles defined as infinitely lethal will cause the goal to be unreachable. Ollis, et al. [Ollis et al., 2008] have addressed this problem by allowing cheap horizontal movement in the goal row of the image space map. This encourages paths to go around the obstacle before entering horizontally, causing the robot to travel around the obstacle as it moves. We modify our force values in a similar fashion. Let  $\mathbf{O}_{n_g, m_g}^\Gamma = \mathbf{O}_g^\Gamma$ . The goal row  $n_g$  of the map  $\mathbf{O}$  is defined to have minimum force value 1.

$$\mathbf{O}_{n_g, 1 \dots w} = 1$$

Figures 5a and 5b display image space paths through identical occupancy grids, except that the latter has been preprocessed to account for obstacle depth.

### 3.5 Navigation

We control the robot based on the relative graph locations of  $\mathbf{O}_s^\Gamma$  and a target node  $\mathbf{O}_{n_{\text{tgt}}, m_{\text{tgt}}}^\Gamma = \mathbf{O}_{\text{tgt}}^\Gamma$  located a predetermined distance  $c_{\text{tgt}}$  along  $P_{\mathbf{O}}^{\min}$ . We find this to be simple and functional; however, Cartesian-based control can alternatively be used by projecting  $\mathbf{O}_{\text{tgt}}^\Gamma$  into robot coordinates. Assuming a center-mounted forward-facing camera,  $\mathbf{O}_s^\Gamma = \mathbf{O}_{h, w/2}^\Gamma$ . In this case, we define forward velocity as a function of the image-space cosine between the column  $w/2$  and the vector from  $\mathbf{O}_s$  to  $\mathbf{O}_{\text{tgt}}$ :

$$speed = \frac{speed_{\max} (h - n_{\text{tgt}})}{\sqrt{(w/2 - m_{\text{tgt}})^2 + (h - n_{\text{tgt}})^2}}$$

where  $speed_{\max}$  is the maximum forward velocity allowed by the robot. Speed increases as  $\mathbf{O}_{\text{tgt}}$  approaches the center of the FOV. In other words, the robot slows during turning maneuvers, which are likely to occur near obstacles. For turns in safe terrain, we increase the speed of the robot by resetting  $speed = speed_{\max}$  if all force values in the bottom half of the center column  $\lfloor w/2 \rfloor$  are less than  $c_{\text{thd}}/3$  after map dilation. Turning angle is defined as the relative horizontal displacement between  $\mathbf{O}_s$  and  $\mathbf{O}_{\text{tgt}}$ , normalized by the FOV  $\theta_w$ .

$$turn = \frac{\theta_w (m_{\text{tgt}} - w/2)}{w}$$

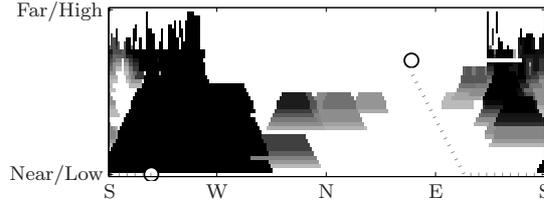


Figure 6: Image space path through the cylindrical occupancy grid. The robot is facing southwest and the goal is east of the robot. It is easier for the robot to turn counter-clockwise, which takes the image space path across the south-south border.

This causes the robot to turn proportionally as the target approaches the edge of the FOV.

## 4 The cylindrical image planner

A shortcoming of the image planner is that the goal can exist outside the FOV. One solution is to use a sensor with a  $2\pi$  radian FOV. In the absence of an omniscam or other panoramic sensor, we can map images to a synthetic panorama over time, based on direction of capture. We call a planner that uses a  $2\pi$  radian panoramic image-space map a *cylindrical planner*. In addition to the assumptions outlined in Section 2.1, our cylindrical planner assumes the existence of global pose information, including orientation. This is necessary to populate and update the cylindrical panorama. One challenge with temporally accumulating panoramas is that information outside the camera’s FOV becomes outdated with movement. In Section 4.2 we describe two techniques for updating data to reflect the robot’s movement.

### 4.1 Panoramic considerations

The cylindrical planner uses a cylindrical occupancy grid  $\mathbf{C}$ .  $\mathbf{C}$  is similar to  $\mathbf{O}$  but contains additional columns to remember information outside the current FOV. Each column in  $\mathbf{C}$  represents a specific compass direction, and data is inserted based on the camera’s orientation at the time of capture.  $\mathbf{C}$  is a radial panorama of what the robot has experienced. Radial panoramas and other forms of mosaics have previously been used for landmark detection and

pose estimation [Chen, 1990, Kelly, 1994, Kelly, 2000, Argyros et al., 2005]. Let  $\varphi$  represent rotation about the vertical axis (yaw) relative to the compass direction North. Information is placed into  $\mathbf{C}$  as a function of  $\varphi$ .

$$\mathbf{C}_{n,m+f(\varphi)} = f(\mathbf{S}_{n,m})$$

Recall that  $f(\mathbf{S}_{n,m})$  is a force metric (Section 3.2). Let  $w_p$  be the number of columns of  $\mathbf{C}$ . We associate columns 1,  $\lfloor w_p/4 \rfloor$ ,  $\lfloor w_p/2 \rfloor$ ,  $\lfloor 3w_p/4 \rfloor$ , and  $w_p$  with the cardinal directions South, West, North, East, and South, respectively (see Figure 6).  $f(\varphi)$ , the mapping from yaw angles to map columns is given by:

$$f(\varphi) = \left\lfloor \frac{w_p(\varphi - \pi)}{2\pi} \bmod w_p \right\rfloor \quad (7)$$

Our robot has two stereo camera rigs that are angled toward the ground. For a virtual cylinder with radius  $r$  aligned with the vertical ( $Z$ ) axis we can identify the intersections of camera pixel rays with the cylinder. For camera position  $\mathbf{p} = [p_x, p_y, p_z]^T$  and pixel ray  $\mathbf{v} = [v_x, v_y, v_z]$  in the robot coordinate frame we want:

$$\begin{aligned} p_x + kv_x &= r \cos \theta, \\ p_y + kv_y &= r \sin \theta. \end{aligned}$$

for scale  $k$ . We square these expressions and add to yield:

$$(p_x^2 + p_y^2) + 2k(p_x v_x + p_y v_y) + k^2(v_x^2 + v_y^2) = r^2.$$

This is a quadratic equation which can be solved for  $k$  in the standard way.

$$\begin{aligned} a &= (v_x^2 + v_y^2) \\ b &= 2(p_x v_x + p_y v_y) \\ c &= (p_x^2 + p_y^2) - r^2 \\ k &= \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \end{aligned}$$

We now determine where each image-pixel ray intersects the cylinder using  $\mathbf{p} + k\mathbf{v}$ . This allows us to identify a range of height and angular values that bound the elements of the cylindrical image. Constructing a tessellation on the cylinder, we project these cylinder

image points to the camera image to yield a mapping between the camera and cylinder images. Values are inserted into the cylinder using this mapping and Equation 7.

Given depth data, another approach is to define a cylinder image origin on the yaw axis, and then reproject reconstructed depth points  $[X_w, Y_w, Z_w]$  to this coordinate frame  $[X_c, Y_c, Z_c]$ .

$$[X_c, Y_c, Z_c] = \frac{1}{\sqrt{X_w^2 + Y_w^2}} [X_w, Y_w, Z_w].$$

Maintaining the 3D locations of points from which cylinder values arise simplifies techniques for propagating and updating cost information as the robot moves.

Let  $\mathbf{C}^\Gamma$  be the 8-connected graph of the occupancy grid  $\mathbf{C}$ . We add arcs between the first and last columns of the occupancy grid graph so that paths can travel across the south-south border of  $\mathbf{C}^\Gamma$ . Arcs exist between  $\mathbf{C}_{n,1}^\Gamma$  and  $\mathbf{C}_{k,w_p}^\Gamma$  for all  $k = \{n-1, n, n+1\}$  and all  $n = \{k-1, k, k+1\}$ , subject to  $1 \leq n \leq h$  and  $1 \leq k \leq h$ . The goal node  $\mathbf{C}_g^\Gamma$  is determined in a similar manner to  $\mathbf{O}_g^\Gamma$ , unless the goal is too close to be projected into the cylinder—in which case we define the goal as achieved. The start node  $\mathbf{C}_s^\Gamma = \mathbf{C}_{h,f(\varphi)}^\Gamma$  is a function of yaw. Figure 6 shows a path from  $\mathbf{C}_s^\Gamma$  to  $\mathbf{C}_g^\Gamma$ , projected into the occupancy grid  $\mathbf{C}$  used to create  $\mathbf{C}^\Gamma$  (the path traverses the south-south border).

As with  $\mathbf{O}$ ,  $\mathbf{C}$  is preprocessed to remove noise and improve path quality. Obstacle width dilation is modified to extend across the south-south border:

$$\mathbf{C}_{n,m} = \max_j (\mathbf{C}_{n,j})$$

where  $j$  is defined as follows:

$$\left. \begin{array}{l} 1 \leq j \leq m + \alpha \\ m - \alpha + w_p \leq j \leq w_p \end{array} \right\} \text{ if } m \leq \alpha$$

$$m - \alpha \leq j \leq m + \alpha \quad \text{if } \alpha < m \leq w_p - \alpha \quad (8)$$

$$\left. \begin{array}{l} 1 \leq j \leq m + \alpha - w_p \\ m - \alpha \leq j \leq w_p \end{array} \right\} \text{ if } w_p - \alpha < m$$

$\alpha$  is defined in Equation 6. The denoising operations are similarly modified. Rotation is

explicitly allowed by resetting values in the bottom row of  $\mathbf{C}$  to the minimum force:

$$\mathbf{C}_{h,1..w} = 1$$

Similarly, finite obstacle depth is accounted for by resetting goal row values to the minimum force if they are within  $\lfloor w_p/4 \rfloor$  columns of  $m_g$ :

$$\mathbf{C}_{n_g,j} = 1$$

where  $j$  is calculated by substituting  $\alpha = \lfloor w_p/4 \rfloor$  and  $m = m_g$  into Equation 8.

Care must be taken to satisfy the constraints of the A\* algorithm. If distance is defined by  $D_i^{\mathbf{S}}$ , the  $L^2$  norm in image space is no longer an admissible heuristic for estimating work because it neglects the possibility of crossing the south-south border. Instead, we use a similar quantity  $L_C^2$  that replaces the horizontal component of the  $L^2$  norm with the shorter of the two horizontal candidates. The  $L_C^2$  norm between  $\mathbf{C}_{n,m}$  and  $\mathbf{C}_{k,j}$  is calculated:

$$L_C^2 = \sqrt{(n-k)^2 + \min((m-j)^2, (w_p - m + j)^2)}$$

$D_i^{\mathbf{R}^f}$  can be used as an admissible heuristic for estimating work because it is found by computing pixel ray intersections with the groundplane model and then computing the Euclidean distances between these points (Section 3.3.2). For the virtual cylinder camera, we compute the groundplane intersections for the cylinder image pixels directly.

## 4.2 Updating techniques

### 4.2.1 Translation based forgetting

Information in  $\mathbf{C}$  is captured when the robot is at a specific location and facing a particular direction. As the robot moves from that location, the data in  $\mathbf{C}$  will cease to be an accurate representation of the associated direction. We propose forgetting data in  $\mathbf{C}$  as a function of translation, and call this technique *translation based forgetting*. We assume that changes in elevation can be ignored to a first approximation and idealize the world as a flat level plane

$\mathbf{R}^f$ . Let  $\mathbf{R}_{\text{bot}}^f$  and  $\tilde{\mathbf{R}}_{\text{bot}}^f$  be the current position of the robot and the position of the robot during the previous update, respectively.  $\tilde{\mathbf{C}}$  and  $\mathbf{C}$  denote the cylinder after the updates at  $\tilde{\mathbf{R}}_{\text{bot}}^f$  and  $\mathbf{R}_{\text{bot}}^f$ , respectively. Let  $\Delta N$ ,  $\Delta E$  be the North and East components of the vector from  $\tilde{\mathbf{R}}_{\text{bot}}^f$  to  $\mathbf{R}_{\text{bot}}^f$ , respectively.  $T$  is the magnitude of the same vector.

$$T = \sqrt{\Delta N^2 + \Delta E^2}$$

The incremental updating function is given by:

$$\mathbf{C}_{n,m} = \max\left(1, \tilde{\mathbf{C}}_{n,m} \frac{c_{\text{fgt}} - T}{c_{\text{fgt}}}\right) \quad (9)$$

where  $c_{\text{fgt}}$  is the translation distance required to erase all information in a single update.  $\mathbf{C}_{n,m}$  cannot decay to less than 1, the minimum force value.  $c_{\text{fgt}}$  is tuned for the desired system behavior—small values of  $c_{\text{fgt}}$  cause the robot to forget information more quickly than large values. The relationship between long-term decay behavior and translation is influenced by the frequency with which Equation 9 is applied.  $\mathbf{C}$  should be updated at a constant rate of time to insure predictable decay vs. distance and velocity.

#### 4.2.2 Depth based updating

If depth data exists for  $\mathbf{S}$ , then  $\mathbf{C}$  can be updated by calculating the motion of depth points relative to the robot. It is convenient to store depth data in a separate ‘cylindrical’ array  $\mathbf{Q}$ . For each force value  $f(\mathbf{S}_{n,m})$  placed into  $\mathbf{C}$  as a function of pixel data  $\mathbf{S}_{n,m}$ , a corresponding 3D position  $\mathbf{R}_{\mathbf{p}}(\mathbf{S}_{n,m})$  in the local robot frame is placed into  $\mathbf{Q}$  at the same location. Let  $\tilde{\mathbf{Q}}$  and  $\mathbf{Q}$  denote the depth cylinder at the robot’s previous location  $\tilde{\mathbf{R}}_{\text{bot}}^f$  and current location  $\mathbf{R}_{\text{bot}}^f$ , respectively. We assume that depth  $\tilde{\mathbf{Q}}_{\tilde{n},\tilde{m}}$  was generated from a point projected to the center of pixel  $\mathbf{S}_{\tilde{n},\tilde{m}}$  when the robot was at  $\tilde{\mathbf{R}}_{\text{bot}}^f$ . Let  $\mathbf{R}_{\mathbf{p}}$  be the point in  $\mathbf{R}$  that generated  $\tilde{\mathbf{Q}}_{\tilde{n},\tilde{m}}$ . As the robot moves, the image projection of  $\mathbf{R}_{\mathbf{p}}$  will change. When 3D points are retained for the cylinder image, we use  $\mathbf{R}_{\text{bot}}^f$  and  $\tilde{\mathbf{R}}_{\text{bot}}^f$  to compute the differential position  $\mathbf{R}_{\Delta}$  between frames. We then apply  $\mathbf{R}_{\Delta}^{-1}$  to the points in  $\tilde{\mathbf{Q}}$  and propagate the points to their new positions in  $\mathbf{Q}$  if they are within the bounds of the cylinder. This provides a

mapping from  $\tilde{\mathbf{Q}}_{\tilde{n},\tilde{m}}$  to  $\mathbf{Q}_{n,m}$ . Information from multiple locations in  $\tilde{\mathbf{Q}}$  may migrate to  $\mathbf{Q}_{n,m}$  as one obstacle is occluded by another.  $\mathbf{Q}_{n,m}$  is updated with the nearest point because closer obstacles present a more immediate navigational hazard. If a non-depth force metric is used to populate  $\mathbf{C}$ , then  $\mathbf{C}$  is updated as follows:

$$\mathbf{C}_{n,m} = \tilde{\mathbf{C}}_{\tilde{n},\tilde{m}}$$

If a depth based force metric is used, then  $\mathbf{C}$  is a synthetic disparity calculated from  $\mathbf{Q}$  by modifying Equation 2 to operate on  $\mathbf{Q}_{n,m}$  instead of  $\mathbf{S}_{n,m}$ :

$$F_i^d = \mathbf{C}_{n,m} = 1 + c_{\text{scl}} \left| \frac{c_f c_b}{\mathbf{Q}_{n,m}} - \frac{c_f c_b}{\mathbf{Q}_{n,m}^f} \right|$$

where  $\mathbf{Q}_{n,m}^f$  is the depth cylinder observed from a flat level plane.

## 5 Hierarchical planners

Image planning provides a variable world view that maximizes map resolution and minimizes the number of map elements required to model the world, given an image sensor reading. However, environmental knowledge is constrained by the current line of sight and far-field data is less refined than near-field data. We believe that the shortcomings of image planning are complimented by the strengths of Cartesian planning and combine the two in a *hybrid hierarchical planner*. The global planner uses a 2D top-down Cartesian occupancy grid  $\mathbf{G}$  to find a coarse path  $P_{\mathbf{G}}^{\min}$ , while the local cylindrical planner navigates toward subgoals  $\mathbf{R}_{\text{sub}}$  a predetermined distance along  $P_{\mathbf{G}}^{\min}$ . To determine what advantages (if any) the local cylindrical planner provides to the hybrid hierarchical planner, we evaluate the proposed system against a double 2D top-down Cartesian hierarchical planner. We refer to the latter system as the *Cartesian hierarchical planner*. Both hierarchical systems use an identical global planner that runs independently on its own processor.

In addition to the assumptions outlined in Section 2.1, the hierarchical planners assume knowledge of global position and orientation. This is necessary to populate, update, and

place start and goal positions in the global map, and to place  $\mathbf{R}_{\text{sub}}$  in the local map. Because  $\mathbf{R}_{\text{sub}}$  is defined to be a specific distance along  $P_{\mathbf{G}}^{\min}$ , any particular  $\mathbf{R}_{\text{sub}}$  will be replaced by a new  $\mathbf{R}_{\text{sub}}$  before the former is actually reached (except when the robot converges on the global goal). This method of planner interaction works well in practice, as long as  $\mathbf{R}_{\text{sub}}$  is within the sensor range of the robot.

## 5.1 The global Cartesian planner

As with the image planner (Section 3) and the cylindrical planner (Section 4), our global Cartesian planner separates the ideas of work, force, and distance. Let  $\mathbf{G}^\Gamma$  represent the 8-connected graph of  $\mathbf{G}$ . The cost of traveling from node  $\mathbf{G}_{k,j}^\Gamma$  to node  $\mathbf{G}_{n,m}^\Gamma$  along arc  $i$  is determined by work  $W_i = F_i D_i$ , where force  $F_i = \mathbf{G}_{n,m}$  and distance  $D_i$  are determined by force and distance metrics, respectively. Map grids are square and we use the  $L^2$  norm (Equation 4) as a distance metric. Values are added to  $\mathbf{G}$  by projecting  $\mathbf{O}$  onto the ground-plane using the robot’s position, orientation, and stereo disparity information. Therefore,  $F_i$  is determined by the force metric used to populate  $\mathbf{O}$ .  $\mathbf{G}$  is preprocessed by a dilation of  $\alpha$ , the map equivalent of a robot half width  $c_{\text{wth}}/2$  plus a buffer  $c_{\text{buf}}$ :

$$\mathbf{O}_{n,m} = \max_{k,j} (\mathbf{O}_{n\pm k, m\pm j})$$

where  $\lfloor -\alpha \rfloor \leq k \leq \lceil \alpha \rceil$ ,  $1 \leq (n \pm k) \leq h$ ,  $\lfloor -\alpha \rfloor \leq j \leq \lceil \alpha \rceil$ , and  $1 \leq (m \pm j) \leq w$ .  $h$  and  $w$  represent the number of rows and columns in  $\mathbf{G}$ , respectively.  $\alpha$  is calculated

$$\alpha = c_{\text{scl}} (c_{\text{wth}}/2 + c_{\text{buf}})$$

where  $c_{\text{scl}}$  is the map scale. The A\* algorithm is used to find the minimum work path  $P_{\mathbf{G}}^{\min}$  from the start node  $\mathbf{G}_{\text{s}}^\Gamma = \mathbf{G}_{n_{\text{s}}, m_{\text{s}}}^\Gamma$  to the goal node  $\mathbf{G}_{\text{g}}^\Gamma = \mathbf{G}_{n_{\text{g}}, m_{\text{g}}}^\Gamma$ , where  $\mathbf{G}_{\text{s}}^\Gamma$  and  $\mathbf{G}_{\text{g}}^\Gamma$  are found by projecting robot and goal coordinates into map space, respectively.  $P_{\mathbf{G}}^{\min}$  is transformed into GPS coordinates  $P_{\mathbf{R}}^{\text{gps}}$  before being passed to a local planner.

## 5.2 The local cylindrical planner

The local cylindrical subsystem of the hybrid hierarchical planner is identical to the system described in Section 4, except that  $\mathbf{C}_g^\Gamma$  is found by projecting  $\mathbf{R}_{\text{sub}}$  into image space, where  $\mathbf{R}_{\text{sub}}$  is the first position along  $P_{\mathbf{R}}^{\text{gps}}$  that is at least  $c_{\text{sub}}$  away from the robot, and  $c_{\text{sub}}$  is a system parameter. Due to the relative frame rates of the local and global planners, the same  $P_{\mathbf{R}}^{\text{gps}}$  may be used to determine multiple  $\mathbf{R}_{\text{sub}}$ .

## 5.3 The local Cartesian planner

The local Cartesian subsystem of the baseline hierarchical planner is similar to the global planner (Section 5.1), except for five differences: (1)  $\mathbf{G}_g^\Gamma$  is found by projecting  $\mathbf{R}_{\text{sub}}$  into local map space, (2) the local path  $P_{\mathbf{G}}^{\text{min}}$  is used for servo control, (3-5) the map has a higher resolution, is fixed in size, and centered on the robot.  $\mathbf{R}_{\text{sub}}$  is defined to be the first position along  $P_{\mathbf{R}}^{\text{gps}}$  that is at least  $c_{\text{sub}}$  from the robot, and the same  $P_{\mathbf{R}}^{\text{gps}}$  may be used to determine multiple  $\mathbf{R}_{\text{sub}}$ . Servoing is accomplished by steering at  $\mathbf{R}_{\text{tgt}}$ , where  $\mathbf{R}_{\text{tgt}}$  is found by projecting  $\mathbf{G}_{\text{tgt}}^\Gamma = \mathbf{G}_{n_{\text{tgt}}, m_{\text{tgt}}}^\Gamma$  into the real world, and  $\mathbf{G}_{\text{tgt}}^\Gamma$  is defined to be  $c_{\text{tgt}}$  along the local path  $P_{\mathbf{G}}^{\text{min}}$ .  $c_{\text{sub}}$  and  $c_{\text{tgt}}$  are a system parameters. Forward velocity and turning angle are defined:

$$\text{speed} = \text{speed}_{\text{max}} \max(0, \cos(\psi_{\mathbf{R}}))$$

$$\text{turn} = \psi_{\mathbf{R}}$$

where  $\psi_{\mathbf{R}}$  is the angular distance between the current heading and  $\mathbf{R}_{\text{sub}}$ , and  $-\pi \leq \psi_{\mathbf{R}} \leq \pi$ . Speed will increase as the target point approaches a position in front of the robot, and turning angle will steer the robot toward the target point.

## 6 The LAGR robot

Our mobile robot platform is provided in conjunction with the DARPA Learning Applied to Ground Robotics (LAGR) program. The length, width, and height dimensions of the robot



Figure 7: The LAGR robot platform.

are approximately 1.2 x .75 x 1.0 meters, respectively. Sensors include: two forward facing Point Grey BumbleBee 2 stereo camera pairs, a Garmin GPS receiver, a magnetic compass, wheel odometers, two forward facing infrared sensors (unused in this work), and a front bumper sensor. The stereo camera pairs output stereo disparity and RGB color data. We have found disparity data to have a maximum range of approximately 15 meters; although the usable range is generally between 5 and 10 meters. There are four processing units: one dedicated to each of the two stereo camera pairs, one for path planning, and one that is a servo controller (the former three computers have dual-core processors). Translation and rotation are achieved via two independently driven front wheels.

## 7 Experiments

All of the methods we have proposed are experimentally evaluated, including: two force metrics ( $F_i^d$  and  $F_i^c$  of Sections 3.2.1 and 3.2.2, respectively), two distance metrics ( $D_i^S$  and  $D_i^{R^f}$  of Sections 3.3.1 and 3.3.2, respectively), and four planning systems (image, cylindrical, hybrid hierarchical, and Cartesian hierarchical from Sections 3, 4, 5, and 5, respectively). We also test two cylindrical updating techniques (translation based forgetting and depth based updating from Sections 4.2.1 and 4.2.2, respectively). Our toy color force metric  $F_i^c$  requires that safe terrain and obstacles be different colors; therefore, we only evaluate  $F_i^c$  in environments where this is the case. As previously stated,  $F_i^c$  is not useful in most environments and is only used to test if our image space systems can function using color.

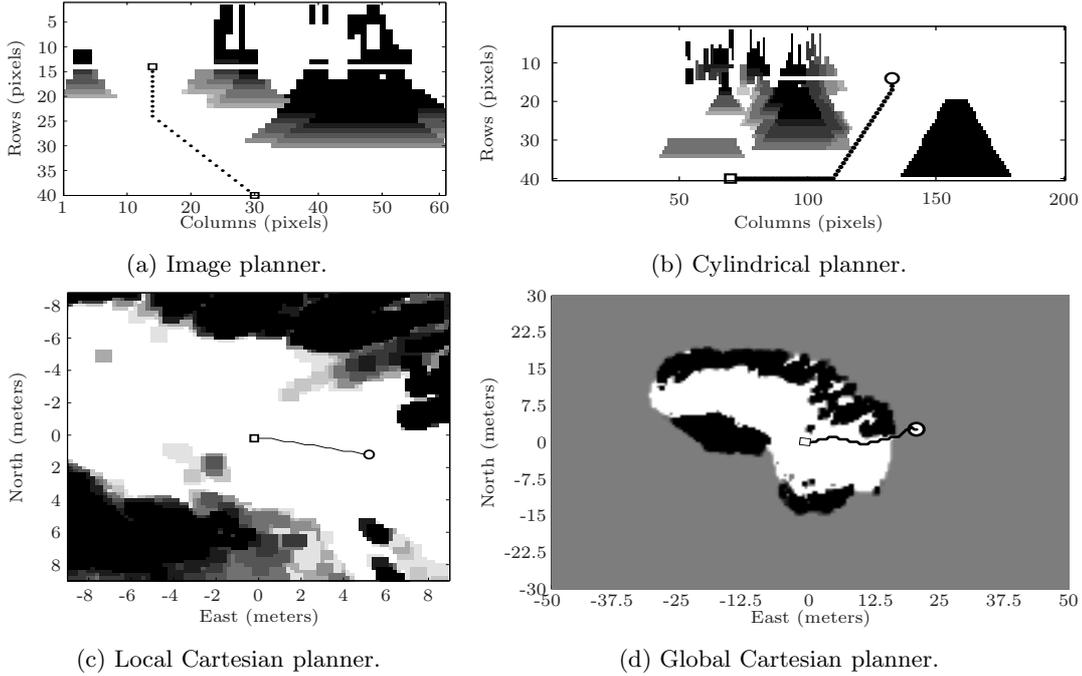


Figure 8: Cost maps with paths, light to dark represents low to high cost, respectively.

The depth based cylindrical updating scheme assumes the existence of depth information; therefore, it is only evaluated in conjunction with the disparity force metric  $F_i^d$ .

We test the following three hypotheses: (1) The flat world distance metric  $D_i^{\mathbf{R}^f}$  will outperform the image distance metric  $D_i^{\mathbf{S}}$  because  $D_i^{\mathbf{R}^f}$  more accurately models real world distances. (2) The cylindrical planner is susceptible to translational oscillation near large obstacles, due to a lack of translational memory. (3) A local image space planner will outperform a local 2D Cartesian planner as a subsystem of a hierarchical planner because the former represents near-field information at a higher resolution.

## 7.1 System parameters

We achieve an effective FOV of 110 degrees (61 grid elements) by having each camera insert information into the cost map (image or cylinder) based on its relative pose. In order to achieve a frame rate of at least 10 hertz, we define the size of the image planner map  $\mathbf{O}$  to be  $h = 40$  grid elements high and  $w = 61$  grid elements wide. Similarly, the cylindrical

cost map  $\mathbf{C}$  is defined to be  $h = 40$  high and  $w_p = 200$  wide (i.e. around the perimeter). In either case, the image from a particular camera is 40 grid elements wide. This resolution is achieved by sub-sampling  $\mathbf{S}$ . Map updates do not incorporate old information. The most recent image determines map values at positions of camera overlap. There is no longer a one-to-one mapping from  $\mathbf{S}$  to  $\mathbf{Q}$  or  $\mathbf{C}$ . However, there is a one-to-one mapping from the subsampled version of  $\mathbf{S}$  to  $\mathbf{Q}$  and  $\mathbf{C}$ . The mask used for denoising the cost map is defined to be 4 pixels high by 1 pixel wide, and is chosen to modify obstacle width as little as possible while still eliminating noise. A cost map from the image planner is displayed in Figure 8a, and a cost map from the cylindrical planner is displayed in Figure 8b.

We set the maximum speed of the robot to 1 meter per second in all experiments. The parameter  $c_{\text{sub}}$  (Section 5.2) used to place the local subgoal  $\mathbf{R}_{\text{sub}}$  is defined to be 5 meters. This distance was chosen to make the local subsystem responsible for navigation around small to medium sized obstacles (e.g. trees, rocks, bushes, and boulders). The image and cylindrical planner parameter  $c_{\text{tgt}}$  (Section 3.5), used to place  $\mathbf{O}_{\text{tgt}}^\Gamma$ , is defined to be the twelfth node along  $P_{\mathbf{O}}^{\min}$ , or the goal node if the path contains fewer than twelve nodes. The parameter  $c_{\text{scl}}$  (Section 3.2) is defined to be 10 or 30 if force is determined by disparity or color, respectively. The parameter  $c_{\text{thd}}$  (Section 3.2) is defined to be 10, and the parameter  $c_{\text{flt}}$  (Section 3.4.1) is defined to be 3. The parameter  $c_{\text{fgt}}$  used in Equation 9 to control translation based forgetting is defined to be 0.4 meters (obstacles are remembered further than 0.4 meters because the frame rate is at least 10 hertz). The values of  $c_{\text{tgt}}$ ,  $c_{\text{scl}}$ ,  $c_{\text{thd}}$ ,  $c_{\text{flt}}$ , and  $c_{\text{fgt}}$  have been tuned by human evaluation of the system’s ability to navigate around a simple obstacle that is 0.5 meters wide, 0.5 meters tall, and 0.25 meters deep.

The local Cartesian planner used in the Cartesian hierarchical system models an 18 by 18 meter square of the world, and the robot is always located in the center of the map graph  $\mathbf{G}^\Gamma$ . This ensures at least 5 meters of the world is modeled between the robot and  $\mathbf{R}_{\text{sub}}$ , and at least 4 meters is modeled between  $\mathbf{R}_{\text{sub}}$  and the map boundary. We believe these distances are suitable for navigation around small to medium sized obstacles. A cost map

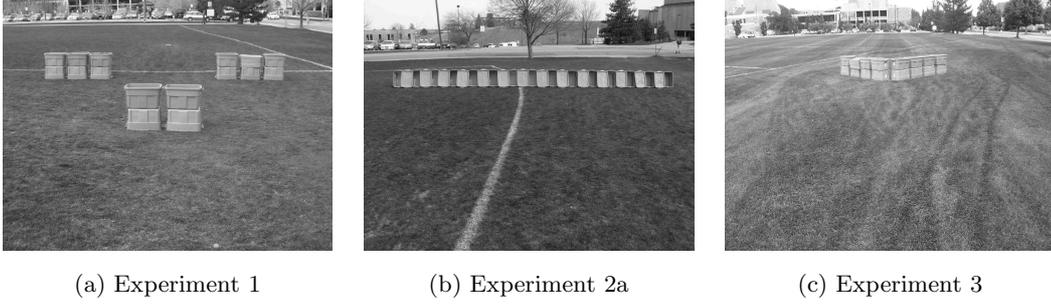


Figure 9: Photos of test courses from the start positions.

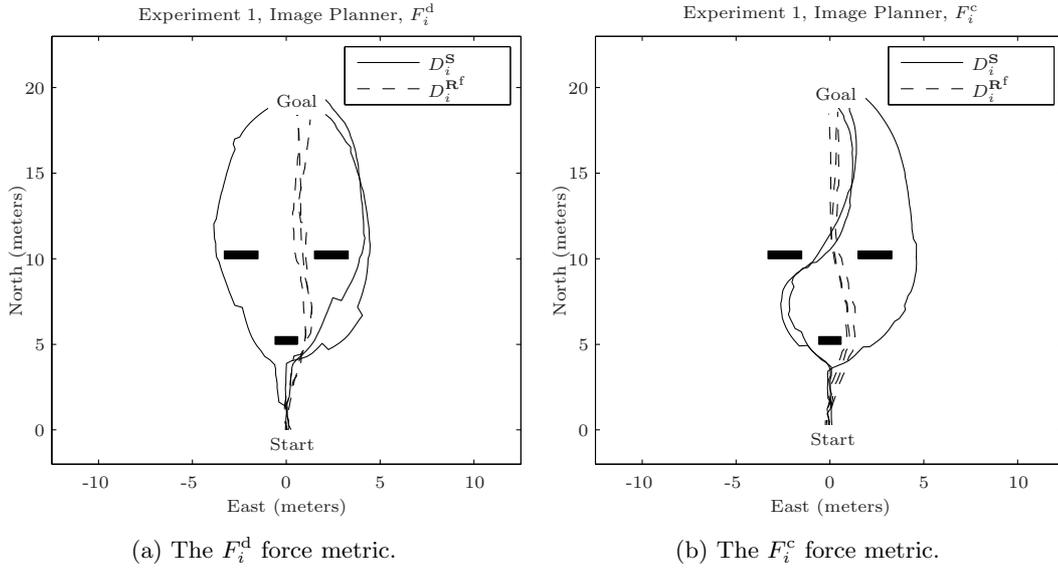


Figure 10: The image planner in Experiment 1, runs using the  $D_i^S$  or  $D_i^R^f$  distance metrics appear solid or dashed, respectively. The axes are in meters.

from the local Cartesian planner is displayed in Figure 8c. In order to make the comparison between the two hierarchical planners as fair as possible, we define the granularity of the local Cartesian occupancy grid  $\mathbf{G}$  to be 0.2 meters. Thus,  $\mathbf{G}^\Gamma$  has 8100 nodes and the local cylindrical map graph  $\mathbf{C}^\Gamma$  has 8000 nodes. The areas of local Cartesian map world elements are 0.04 meters squared. The areas of image cylindrical map world elements vary from 0.019 to  $\infty$  meters squared, from the near- to far-field, respectively. The frame rate of the local Cartesian planner is comparable to that of the local cylindrical planner.

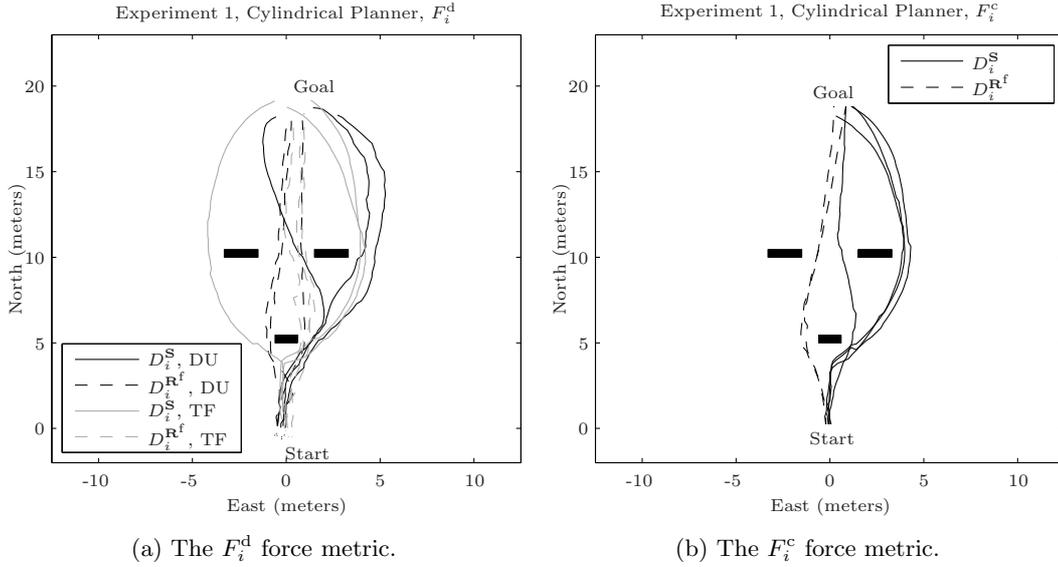


Figure 11: The cylindrical planner in Experiment 1. Runs using the  $D_i^S$  or  $D_i^{R^f}$  distance metrics appear solid or dashed, respectively. Runs using depth based updating or translation based forgetting are black or gray, respectively. The axes are in meters.

## 7.2 Experiment 1: shakeout course

Experiment 1 consists of the NIST (National Institute of Standards and Technology) LAGR shakeout course (Figure 9a). The goal is placed 20 meters north of the starting location. The course contains three obstacles constructed out of 4, 6, and 6 large plastic tubs, respectively. The first obstacle is placed 5 meters toward the goal, directly in front of the robot. The next two obstacles are placed 5 meters further, separated by 3 meters. The width of the obstacles are 1.2, 1.8, and 1.8 meters, respectively, and the length and height of the obstacles are 0.4 and 0.8 meters, respectively. The planning systems tested in Experiment 1 include all combinations of the image and cylindrical planners,  $F_i^d$  and  $F_i^c$  force metrics, and  $D_i^S$  and  $D_i^{R^f}$  distance metrics. Additionally, both cylindrical updating schemes (translation based forgetting and depth based updating) are tested in conjunction with a  $F_i^d$  cylindrical planner. Each combination is run 3 times, resulting in 30 total runs.

The primary purpose of Experiment 1 is to show that image planning is possible over a wide range of system implementations. Additionally, we hope to observe some discriminating

Table 1: Experiment 1 runtimes in seconds

Image Planner				
System	Runs 1-3			Mean
$F_i^d D_i^S$	30.60	31.07	30.26	30.64
$F_i^d D_i^{R^f}$	24.89	25.17	25.28	25.11
$F_i^c D_i^S$	34.06	33.00	32.02	33.03
$F_i^c D_i^{R^f}$	24.94	25.16	25.11	25.07
Cylindrical Planner				
System	Runs 1-3			Mean
DU $F_i^d D_i^S$	30.66	34.43	33.15	32.75
DU $F_i^d D_i^{R^f}$	25.25	28.62	25.02	26.30
TF $F_i^d D_i^S$	30.75	30.54	31.87	31.05
TF $F_i^d D_i^{R^f}$	24.19	26.02	25.09	25.10
TF $F_i^c D_i^S$	31.14	30.25	29.40	30.26
TF $F_i^c D_i^{R^f}$	25.64	25.22	25.43	25.43

Table 2: Experiment 2a runtimes in seconds

Image Planner				
System	Runs 1-3			Mean
$F_i^d D_i^S$	$\infty$	$\infty$	$\infty$	$\infty$
$F_i^d D_i^{R^f}$	$\infty$	$\infty$	$\infty$	$\infty$
$F_i^c D_i^S$	35.31	$\infty$	$\infty$	$\infty$
$F_i^c D_i^{R^f}$	$\infty$	$\infty$	$\infty$	$\infty$
Cylindrical Planner				
System	Runs 1-3			Mean
DU $F_i^d D_i^S$	41.72	46.37	44.94	44.34
DU $F_i^d D_i^{R^f}$	$\infty$	$\infty$	53.92	$\infty$
TF $F_i^d D_i^S$	39.52	$\infty$	49.07	$\infty$
TF $F_i^d D_i^{R^f}$	65.84	$\infty$	46.51	$\infty$
TF $F_i^c D_i^S$	39.46	40.42	49.22	43.03
TF $F_i^c D_i^{R^f}$	$\infty$	69.26	$\infty$	$\infty$

trends along the different dimensions of our system space. GPS data from each run is superimposed on an illustration of the course in Figures 10 and 11 and runtimes are displayed in Table 1 (note that TF and DU denote translation based forgetting and depth based updating, respectively). Image planning systems are displayed in Figure 10, while cylindrical planning systems appear in Figure 11. Figures 10a and 11a display tests using the  $F_i^d$  force metric, while Figures 10b and 11b show tests using the  $F_i^c$  force metric. The  $D_i^S$  and  $D_i^{R^f}$  distance metrics are plotted with solid and dashed lines, respectively.

### 7.3 Experiment 2: walls of 10 and 50 meters

Experiment 2a consists of a single wall placed 10 meters east of the starting location. The goal is positioned another 10 meters east of the wall. The width, length, and height dimensions

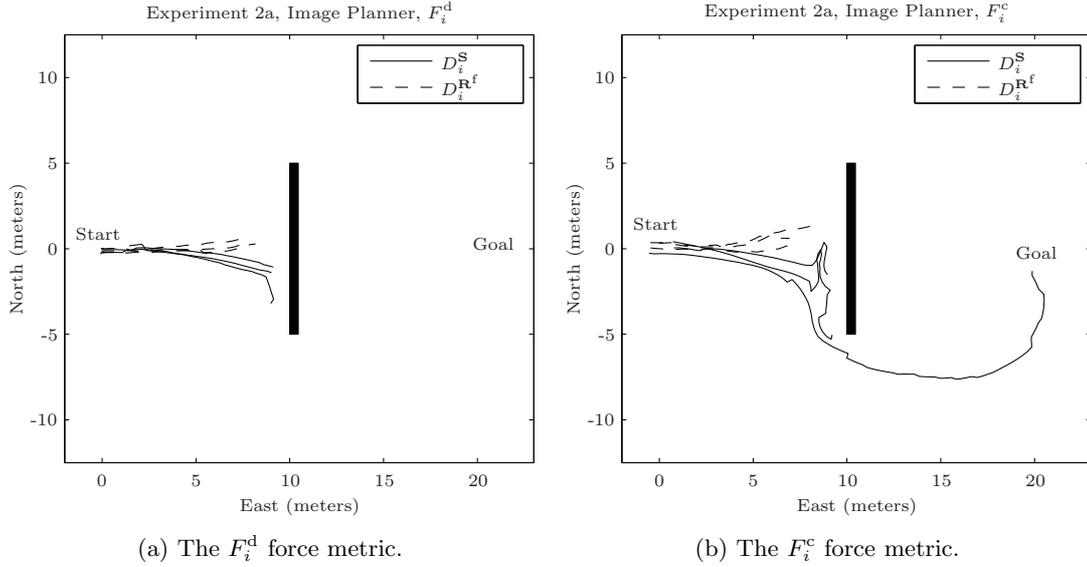


Figure 12: The image planner in Experiment 2a, runs using the  $D_i^S$  or  $D_i^{R^f}$  distance metrics appear solid or dashed, respectively. The axes are in meters.

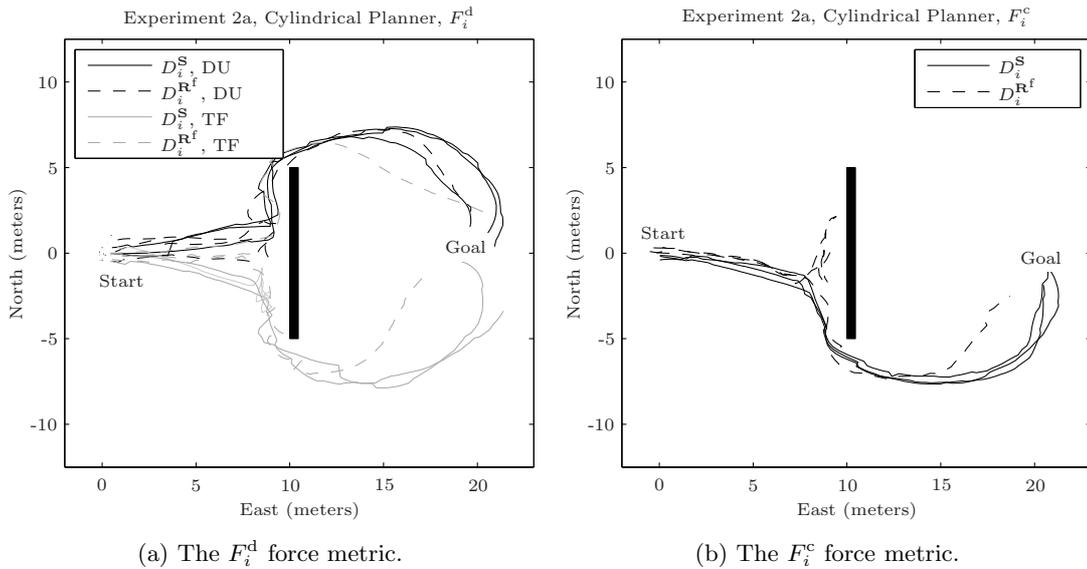


Figure 13: The cylindrical planner in Experiment 2a. Runs using the  $D_i^S$  or  $D_i^{R^f}$  distance metrics appear solid or dashed, respectively. Runs using depth based updating or translation based forgetting are black or gray, respectively. The axes are in meters.

of the wall are 10, 0.4, and 0.6 meters, respectively. A photo of the course is displayed in Figure 9b. The course consists of the aforementioned blue tubs and grassy field. The same planning systems are tested as in Experiment 1, and trials are repeated 3 times. The actual

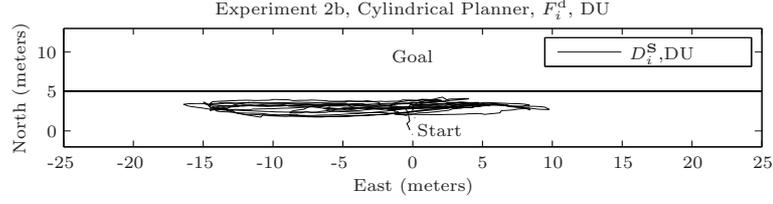


Figure 14: Experiment 2b, the  $F_i^d$  cylindrical planner using  $D_i^S$  and depth based updating.

Table 3: Experiment 3 runtimes in seconds

Image Planner				
System	Runs 1-3			Mean
$F_i^d D_i^S$	20.53	33.01	56.42	36.65
$F_i^d D_i^{R^f}$	57.98	22.76	38.62	39.79
$F_i^c D_i^S$	23.11	22.80	22.07	22.66
$F_i^c D_i^{R^f}$	21.69	23.14	22.13	22.32
Cylindrical Planner				
System	Runs 1-3			Mean
DU $F_i^d D_i^S$	19.82	32.91	22.24	24.99
DU $F_i^d D_i^{R^f}$	26.35	32.68	22.81	27.28
TF $F_i^d D_i^S$	$\infty$	68.15	38.18	$\infty$
TF $F_i^d D_i^{R^f}$	21.74	74.90	49.85	48.83
TF $F_i^c D_i^S$	24.17	23.06	24.63	23.96
TF $F_i^c D_i^{R^f}$	21.46	32.47	32.78	28.90

Table 4: Experiment 2b runtimes in seconds

Cylindrical Planner				
System	Runs 1-3			Mean
DU $F_i^d D_i^S$	$\infty$	$\infty$	$\infty$	$\infty$

Table 5: Experiment 4 runtimes in seconds

Hybrid Hierarchical Planner			
System	Runs 1-2		Mean
DU $F_i^d D_i^S$	122.4	116.7	119.55
DU $F_i^d D_i^{R^f}$	104.2	106.1	105.15
TF $F_i^d D_i^S$	120.9	132.4	126.65
TF $F_i^d D_i^{R^f}$	102.3	103.6	102.95
Cartesian Hierarchical Planner			
System	Runs 1-2		Mean
Cartesian	140.9	121.4	131.15

GPS data from each run is superimposed on an illustration of the course in Figures 12 and 13. The planning systems displayed in Figures 12a, 12b, 13a, and 13b are identical to those displayed in Figures 10a, 10b, 11a, and 11b, respectively. Runtimes are displayed in Table 2.

Experiment 2a is designed to test our hypothesis that the cylindrical planner is vulnerable to translational oscillation when large obstacles obscure the goal. Additionally, we expect to observe rotational oscillation when the goal is blocked by obstacles wider than the FOV.

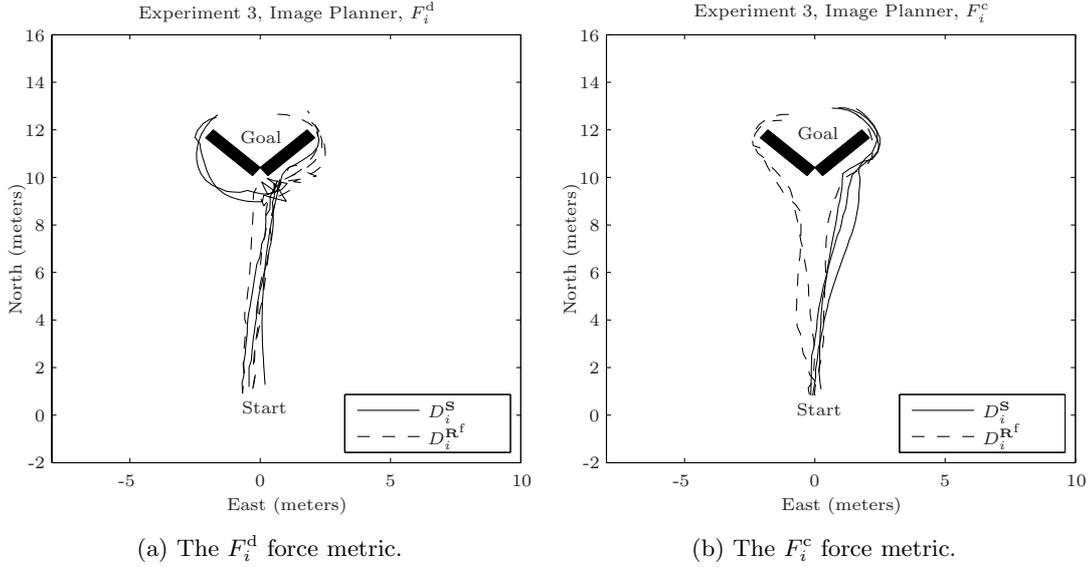


Figure 15: The image planner in Experiment 3, runs using the  $D_i^S$  or  $D_i^{R^f}$  distance metrics appear solid or dashed, respectively. The axes are in meters.

A test is halted if the robot oscillates 8 times without making any forward progress toward the goal (the oscillation count resets to 0 whenever progress is made). Elapsed time values of  $\infty$  denote that a system did not complete the course.

Experiment 2b consists of a single wall placed 5 meters north of the starting location. The goal is positioned another 5 meters north of wall. The width, length, and height dimensions of the wall are 50, 0.1, and 1 meters, respectively (Figure 14). We test the cylindrical planner using the force-from-disparity metric  $F_i^d$  in conjunction with  $D_i^S$  distance and depth based updating. The system is run 3 times. Experiment 2b is designed to test the limits of this particular planner because it completed Experiment 2a all three times.

#### 7.4 Experiment 3: V shaped course

Experiment 3 consists of a symmetrical ‘V’ shaped obstacle placed 10 meters north of the starting location. The ‘V’ consists of two 2.4 meter long walls joined at an angle of 97 degrees. The depth and height of the obstacles are 0.4 and 0.8 meters, respectively. The goal is centered between the far ends of the two walls. A photo of the course is displayed

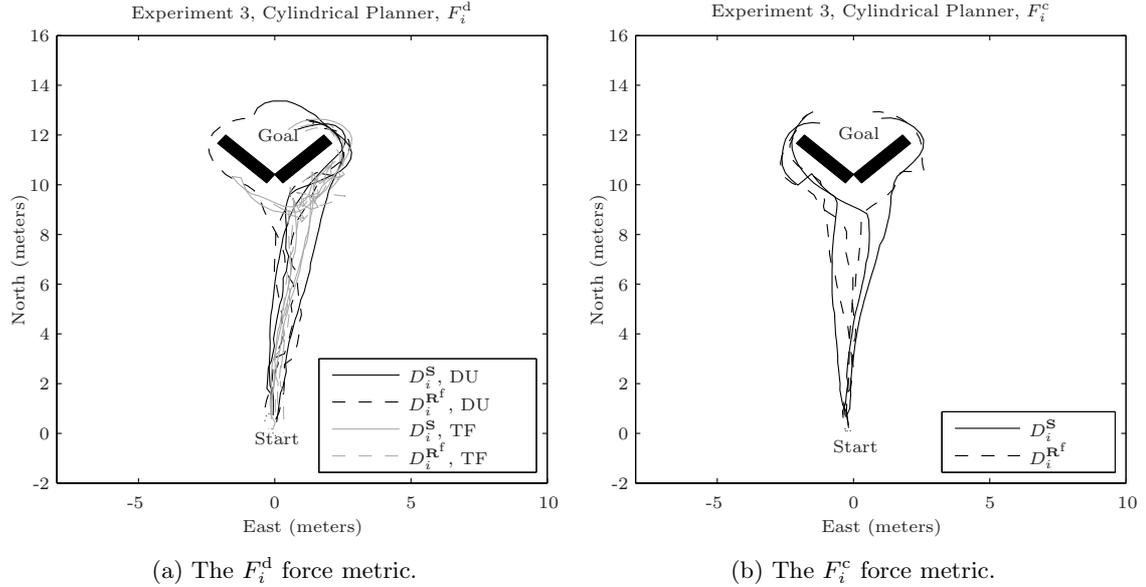
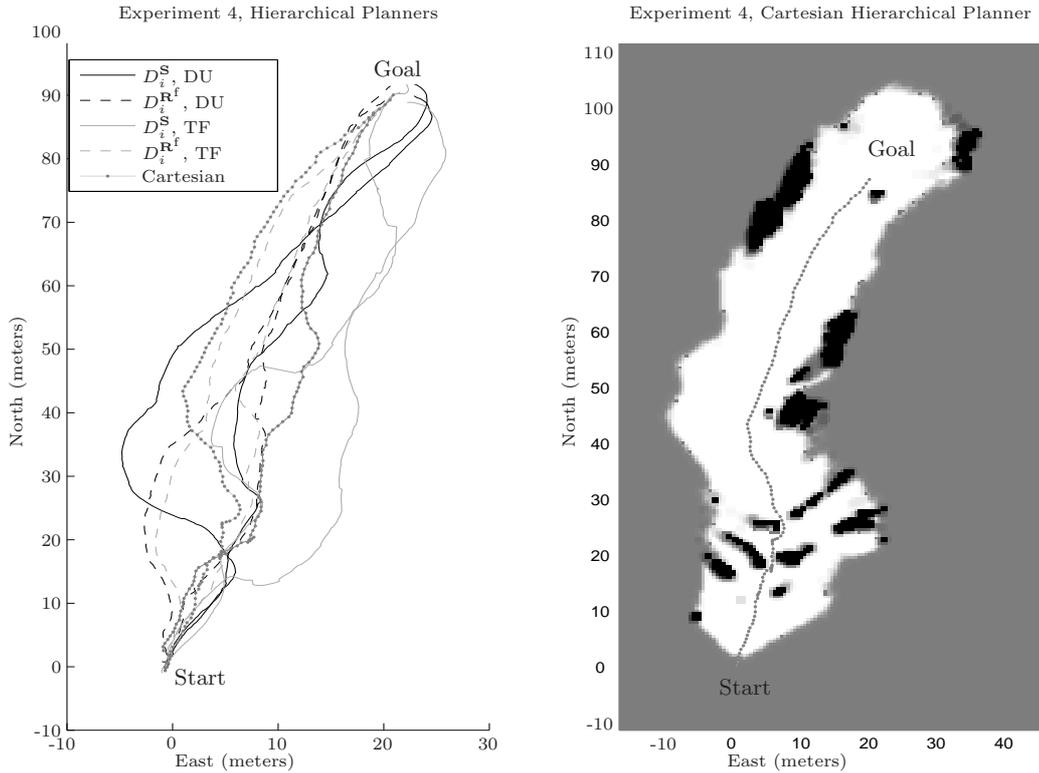


Figure 16: The cylindrical planner in Experiment 3. Runs using the  $D_i^S$  or  $D_i^{R^f}$  distance metrics appear solid or dashed, respectively. Runs using depth based updating or translation based forgetting are black or gray, respectively. The axes are in meters.

in Figure 9c. The course uses the same blue tubs and grassy field as Experiments 1 and 2a, and the same planning systems are also tested. Trials are repeated 3 times. The actual GPS data from each run is superimposed on an illustration of the course in Figures 15 and 16. The planning systems displayed in Figures 15a, 15b, 16a, and 16b are identical to those displayed in Figures 10a, 10b, 11a, and 11b. Experiment 3 evaluates the image and cylindrical planners on a controlled course not specifically designed to induce pathological behavior. The runtimes are given in Table 3.

## 7.5 Experiments 4-7: natural terrain courses

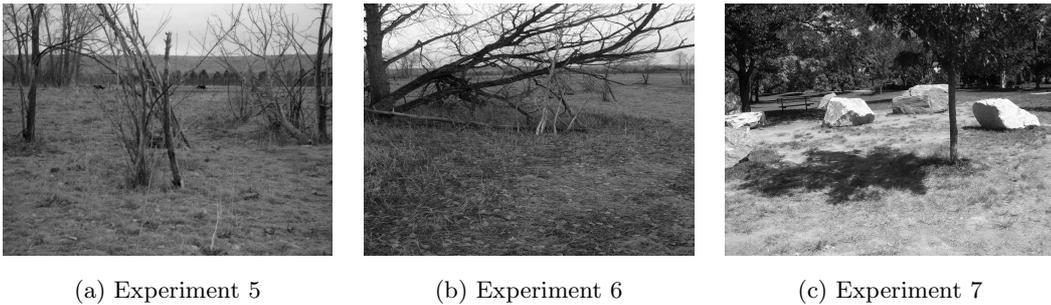
Experiments 4-7 are in natural terrain. They are conducted to demonstrate that the hybrid and Cartesian hierarchical systems actually perform in unknown unstructured outdoor environments, and also to determine if the behavior of the stand-alone cylindrical planning systems transfers to hybrid hierarchical systems. Experiment 4 consists of boulders and trees on a lightly undulating ground surface with grass. The course is 95 meters long. GPS paths



(a) Hybrid hierarchical runs using  $D_i^S$  or  $D_i^{R^f}$  are solid or dashed, respectively, runs using depth based updating or translation based forgetting are black or gray, respectively, and all use  $F_i^d$ . Runs from the Cartesian hierarchical planner are dotted gray.

(b) The Cartesian hierarchical planner's first run superimposed on the corresponding global occupancy grid. Low to high force values are displayed in white to black, respectively.

Figure 17: Experiment 4 GPS paths. The axes are in meters.



(a) Experiment 5

(b) Experiment 6

(c) Experiment 7

Figure 18: Photos of Experiments 5, 6, and 7 from the goal, start, and start, respectively.

are displayed in Figure 17a. An approximate representation of the course can be seen in Figure 17b. All systems are tested twice. The runtimes are displayed in Table 5.

Experiments 5 and 6 contain shrubs, trees, and dead grass ranging in height from 3 centime-

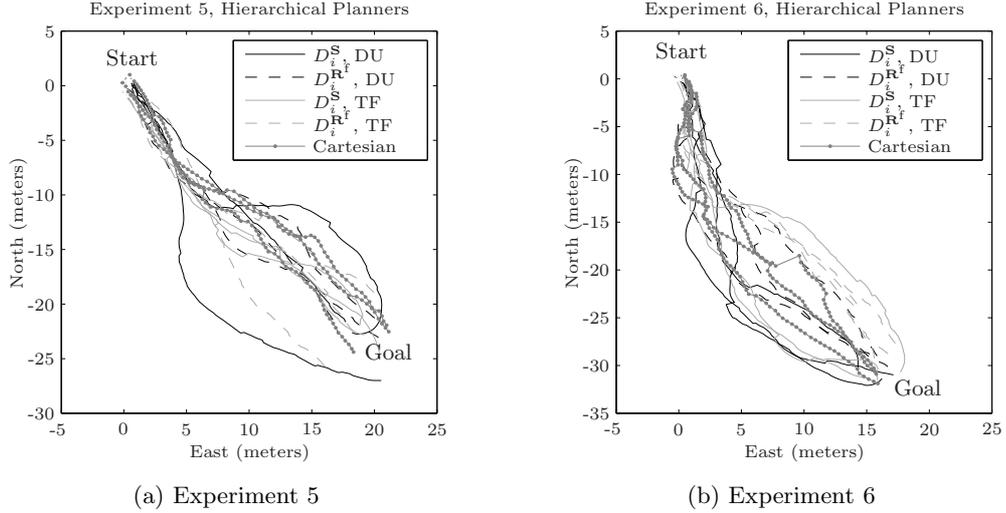


Figure 19: Experiment 5 and 6, hybrid hierarchical runs using  $D_i^S$  or  $D_i^{R^f}$  are solid or dashed, respectively, runs using depth based updating or translation based forgetting are black or gray, respectively, and all use  $F_i^d$ . Runs from the Cartesian hierarchical planner are dotted gray. The axes are in meters.

Table 6: Experiment 5 runtimes in seconds

Hybrid Hierarchical Planner				
System	Runs 1-3			Mean
DU $F_i^d$ $D_i^S$	35.99	48.54	47.03	43.85
DU $F_i^d$ $D_i^{R^f}$	37.81	39.99	39.55	39.12
TF $F_i^d$ $D_i^S$	63.36	41.05	36.78	47.06
TF $F_i^d$ $D_i^{R^f}$	37.38	35.48	36.26	36.37
Cartesian Hierarchical Planner				
System	Runs 1-3			Mean
Cartesian	40.40	41.56	39.61	40.52

Table 7: Experiment 6 runtimes in seconds

Hybrid Hierarchical Planner				
System	Runs 1-3			Mean
DU $F_i^d$ $D_i^S$	53.64	48.65	43.96	48.75
DU $F_i^d$ $D_i^{R^f}$	53.24	46.69	42.76	47.56
TF $F_i^d$ $D_i^S$	44.39	45.15	48.91	46.15
TF $F_i^d$ $D_i^{R^f}$	41.56	44.10	45.46	43.71
Cartesian Hierarchical Planner				
System	Runs 1-3			Mean
Cartesian	47.77	48.36	43.64	46.59

ters to 0.5 meters. The terrain is rocky, with rocks protruding above the ground surface 4 to 12 centimeters. Photos of the courses are displayed in Figures 18a and 18b, respectively. In Experiment 6 there is a 10 meter wide stand of interwoven trees blocking the robot's view of the goal from the starting position. The courses are 30 and 35 meters long, respectively.

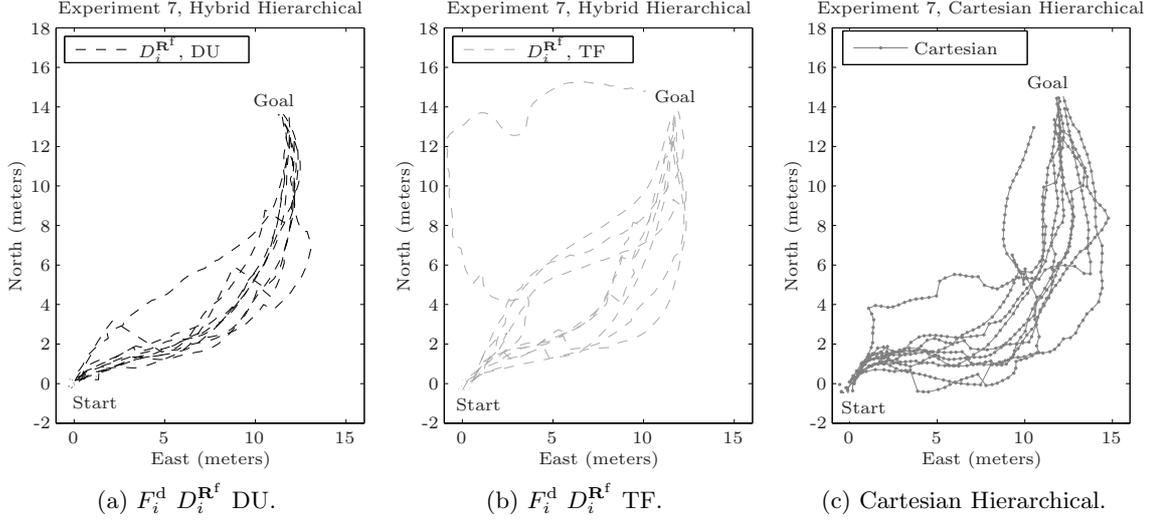


Figure 20: GPS paths of the hybrid and Cartesian hierarchical planners in Experiment 7.

GPS paths are displayed in Figures 19a and 19b, respectively. All systems are tested three times, and the runtimes are given in Tables 6 and 7, respectively.

Experiment 7 evaluates the hierarchical planners on a grass and dirt field with trees and boulders as obstacles. A photo of the course is displayed in Figure 18c. The course is 19 meters long. The purpose of Experiment 7 is to compare the best performing hybrid hierarchical planners (as determined by previous experiments) against the Cartesian hierarchical planner with enough trials to provide statistical significance. Hybrid hierarchical planners using the  $D_i^{\mathbf{R}^f}$  distance metric with either cylindrical updating strategy are tested. We selected this course because we believe cluttered environments give hybrid hierarchical planners an advantage over the Cartesian hierarchical planner. All systems are tested 10 times. GPS paths are displayed in Figure 20, and the runtimes are given in Table 8.

## 8 Discussion

### 8.1 Force metrics: $F_i^{\mathbf{d}}$ , $F_i^{\mathbf{c}}$

We use force metrics from disparity  $F_i^{\mathbf{d}}$  and color  $F_i^{\mathbf{c}}$  to evaluate the effects of different sensors and force definitions on image planning. Comparison between  $F_i^{\mathbf{d}}$  and  $F_i^{\mathbf{c}}$  is only

Table 8: Experiment 7 runtimes in seconds

Hybrid Hierarchical Planner							
System	Runs 1-10					Mean	Std.
	27.41	26.35	25.59	29.78	32.47		
DU $F_i^d$ $D_i^{R^f}$	31.40	27.38	28.84	28.82	26.59	28.46	2.24
	24.40	27.33	35.89	25.59	42.84		
TF $F_i^d$ $D_i^{R^f}$	44.58	35.69	25.96	26.07	25.53	31.39	7.71
Cartesian Hierarchical Planner							
System	Runs 1-10					Mean	Std.
	37.38	30.45	40.83	30.53	33.15		
Cartesian	34.78	42.55	51.09	34.23	57.64	39.26	9.01

possible on the subset of courses for which  $F_i^c$  is tested: Experiments 1, 2a, and 3 (the NIST shakeout, 10 meter wall, and V shaped courses, respectively). In Experiment 1, both  $F_i^d$  and  $F_i^c$  are effective and there is no clear winner with respect to runtime. In contrast, Experiments 2a and 3 show  $F_i^c$  systems encountering oscillation less frequently and progressing further when oscillation does occur. This suggests  $F_i^c$  has an advantage over  $F_i^d$ , which we attribute to the greater range of color vs. disparity. Systems using  $F_i^c$  are more likely to see and avoid an obstacle before moving close enough to induce oscillation. It is worth restating that any *predefined* color force metric is only useful in a subset of environments. The obstacle and ground colors of Experiments 1, 2a, and 3 were chosen to be compatible with  $F_i^c$ ; therefore, the superiority of  $F_i^c$  to  $F_i^d$  may be due to course appearance.

## 8.2 Distance metrics: $D_i^S$ , $D_i^{R^f}$

$D_i^S$  and  $D_i^{R^f}$  define distance as the  $L^2$  norm in image space and its projected distance along a flat level groundplane, respectively. We have hypothesized that  $D_i^{R^f}$  will outperform  $D_i^S$  because the former better approximates real world distances than the latter. Our experiments

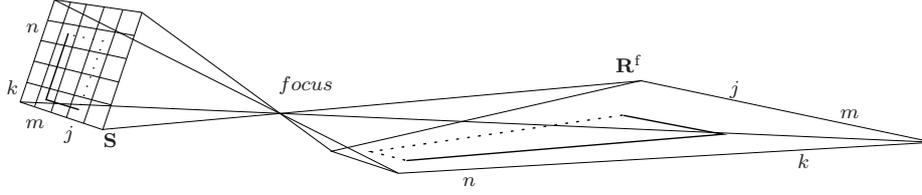


Figure 21: A projection of two paths from image space  $\mathbf{S}$  onto a flat groundplane  $\mathbf{R}^f$ . The distance along  $\mathbf{R}^f$  is used for the  $D_i^{\mathbf{R}^f}$  metric. Although both paths have the same number of vertical and horizontal edges, the dotted path is shorter due to perspective. The solid and dashed paths appear reversed on the image plane due to projection through the focus.

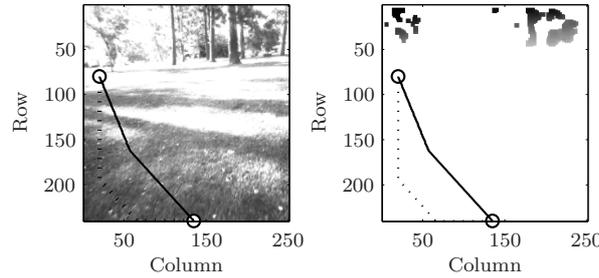


Figure 22: Image space paths using the  $D_i^{\mathbf{S}}$  and  $D_i^{\mathbf{R}^f}$  distance metrics, solid and dotted, respectively, through a scene image (left) and corresponding occupancy grid (right). In an empty field,  $D_i^{\mathbf{S}}$  causes the path to change direction further away from the robot than  $D_i^{\mathbf{R}^f}$ .

validate this hypothesis for cases without rotational or translational oscillation. However, they also show that  $D_i^{\mathbf{S}}$  is less susceptible to oscillation than  $D_i^{\mathbf{R}^f}$ . In experiments free from oscillation, systems using  $D_i^{\mathbf{R}^f}$  achieve the goal faster than those using  $D_i^{\mathbf{S}}$ . Examples include the NIST shakeout and natural terrain courses (Experiments 1, 4-6, respectively). On the other hand, in cases where terminal oscillation is observed,  $D_i^{\mathbf{S}}$  explores more terrain despite the oscillation (e.g. the Experiment 2 wall courses). Thus,  $D_i^{\mathbf{S}}$  sometimes allows a system to discover a transition out of an oscillatory state and achieve the goal when  $D_i^{\mathbf{R}^f}$  cannot. This exploratory behavior also means that  $D_i^{\mathbf{S}}$  risks wasting fuel and time if such a transition is never found. The observed trends extend to the V shaped course (Experiment 3), with the exception of the  $D_i^{\mathbf{R}^f}$  translation-based-forgetting cylindrical planner. In the latter case, one  $D_i^{\mathbf{S}}$  run is terminated after oscillating 8 times, while similar behavior in the  $D_i^{\mathbf{R}^f}$  system

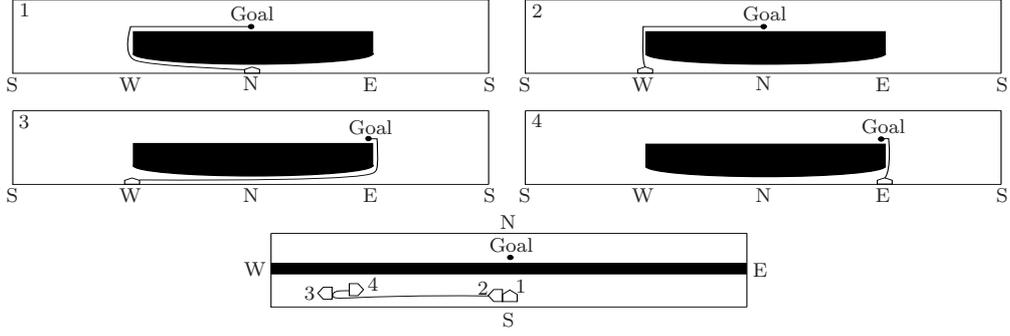


Figure 23: Translational oscillation caused by a large obstacle. Panes 1-4 show the paths through the cylinder that are generated at the respective locations in the bottom diagram. (1) The robot decides to move west. (2) Directional change reinforces the decision to go west. (2-3) Westward translation causes the goal to move east, relative to the robot. (3) It becomes cheaper to go east. (4) Directional change reinforces the decision to go east.

resolves itself before 8 oscillations. An interesting difference between the two distance metrics is that, in the absence of obstacles,  $D_i^{\mathbf{R}^f}$  induces the robot to move directly toward the goal, while  $D_i^{\mathbf{S}}$  tends to produce navigation along ‘hook’ shaped trajectories. This behavior can be observed in all experiments. The extra distance traveled by systems using  $D_i^{\mathbf{S}}$  is the primary reason they tend to achieve the goal more slowly than systems using  $D_i^{\mathbf{R}^f}$ .

All of the observed differences are explained by the characteristics of the two metrics.  $D_i^{\mathbf{S}}$  defines the distance between vertically or horizontally neighboring nodes as 1 and diagonally neighboring nodes as  $\sqrt{2}$ . The combined distance required to move horizontally from  $\mathbf{O}_{n,m}^{\Gamma}$  to  $\mathbf{O}_{n,j}^{\Gamma}$  then vertically to  $\mathbf{O}_{k,j}^{\Gamma}$  is equal to the distance required to move vertically from  $\mathbf{O}_{n,m}^{\Gamma}$  to  $\mathbf{O}_{k,m}^{\Gamma}$  then horizontally to  $\mathbf{O}_{k,j}^{\Gamma}$ . In contrast,  $D_i^{\mathbf{R}^f}$  represents distance as the projected length of ground-surface between pixels centers in  $\mathbf{S}$ , and the distance required to move between two vertically, horizontally, or diagonally neighboring nodes increases as grid row  $n$  approaches the horizon. Thus, if  $n > k$  (i.e.  $n$  is closer to bottom of the FOV than  $k$ ), then the combined distance required to move horizontally from  $\mathbf{O}_{n,m}^{\Gamma}$  to  $\mathbf{O}_{n,j}^{\Gamma}$  then vertically to  $\mathbf{O}_{k,j}^{\Gamma}$  will be less than the combined distance required to move vertically from  $\mathbf{O}_{n,m}^{\Gamma}$  to  $\mathbf{O}_{k,m}^{\Gamma}$  then horizontally to  $\mathbf{O}_{k,j}^{\Gamma}$  (see Figure 21). This is because if  $n > k$ , then the distance

between  $\mathbf{O}_{n,m}^\Gamma$  and  $\mathbf{O}_{n,j}^\Gamma$  is less than the distance between  $\mathbf{O}_{k,m}^\Gamma$  and  $\mathbf{O}_{k,j}^\Gamma$ .

Given constant force values,  $D_i^{\mathbf{R}^f}$  will cause a path to perform as much horizontal movement in the bottom of the FOV (i.e. the near-field) as possible. Thus,  $D_i^{\mathbf{R}^f}$  encourages a system to steer toward the goal immediately, while  $D_i^{\mathbf{S}}$  allows turning movement to be postponed until later. This is illustrated in Figure 22. The  $D_i^{\mathbf{S}}$  ‘hook’ shaped real-world robot trajectories (shown in Section 7) are a result of this belated steering behavior, as is the propensity for increased exploration in situations that cause oscillation.

$D_i^{\mathbf{R}^f}$  is relatively susceptible to translational oscillation because arc distance in  $\mathbf{O}^\Gamma$  increases as a function of image height. Small changes in far-field information, including the relative position of the goal, can trigger large corrections in near-field portions of the path. This is observed on the 10 meter wall course (Experiment 2a, Figure 13). Consider the case where a large obstacle exists between the robot and the goal (Figure 23). The initial rotation away from the obstacle reinforces the decision to move in a particular direction by decreasing the optimal path length. As the robot moves, the relative location of  $\mathbf{C}_g^\Gamma$  moves in the opposite direction of travel. Eventually,  $D_i^{\mathbf{R}^f}$  induces a near-field correction and the robot changes direction. Assuming a unit force associated with non-obstacle portions of  $\mathbf{C}$ , this is guaranteed to repeat whenever the distance added to the path in the far-field (due to the relative movement of the goal vs. the robot) is greater than the near-field distance required to change direction. Although,  $D_i^{\mathbf{S}}$  is susceptible to oscillation given a large enough obstacle (e.g. the 50 meter wall in Experiment 2b), it allows movement to continue further along the obstacle before a direction reversal occurs. This is because near-field movement involves the same amount of distance as far-field movement, and a simple rotation can substantially change the length of a path. The apparent location of  $\mathbf{C}_g^\Gamma$  must approach the opposite edge of the obstacle before a direction reversal occurs.

### 8.3 Cylindrical updating: translation based forgetting, depth based updating

Translation based forgetting decreases the force of an obstacle as a function of robotic displacement, while depth based updating explicitly tracks obstacle movement relative to the robot. The former is tested in conjunction with all cylindrical and hybrid planners, while the latter is only tested with systems using the  $F_i^d$  force metric. Translation based forgetting slightly outperforms depth updating in situations without translational oscillation, but depth updating is more robust to translational oscillation. On the 10 meter wall course (Experiment 2a, Figure 13a) the  $F_i^d D_i^S$  cylindrical planner experiences translational oscillation and systems using depth based updating always escape, while systems using translation based forgetting only escape twice. These trends extend to the V shaped course (Experiment 3, Figure 16a, Table 3), where translation based forgetting experiences more severe oscillation than depth based updating.

With translation based forgetting, the robot will always forget an obstacle of a particular force  $\mathbf{O}_{n,m}$  after moving a predetermined distance  $c_{\text{dst}}$  at a constant speed  $speed$ .  $c_{\text{dst}}$  is calculated by recursively expanding Equation 9, assuming it is used at a constant rate  $c_{\text{rt}}$ :

$$1 \leq \mathbf{O}_{n,m} \prod_{u=1}^{\lceil \frac{c_{\text{dst}} c_{\text{rt}}}{speed} \rceil} \frac{c_{\text{fgt}} - \frac{speed}{c_{\text{rt}}}}{c_{\text{fgt}}}$$

where  $speed/c_{\text{rt}}$  is the distance traveled between calls to Equation 9. The oscillation amplitude can be tuned by adjusting  $c_{\text{fgt}}$ ,  $c_{\text{rt}}$ , and  $speed$ ; however, it will be constant given a parameter set. If the robot encounters an obstacle longer than this distance, then the robot will turn back toward the goal before it has circumvented the obstacle. Oscillation occurs because the robot must turn around before rediscovering the forgotten portion of the obstacle. After the turn, it requires less graph distance to reach the goal by traveling in the new direction—especially in the case of  $D_i^S$ . Depending on the specific system parameters and apparent location of  $\mathbf{C}_g^\Gamma$ , translational oscillation may be induced by the distance metric before it is caused by translation based forgetting. Force values are stored at constant headings in the cylinder; therefore, after turning away from an obstacle, the robot will avoid turning

back toward its original heading until the obstacle is forgotten. Large values of  $c_{\text{fgt}}$  and slow frame-rates tend to cause the robot to over-avoid small obstacles (i.e. move unnecessarily far around them). The forgetting parameters must be tuned to elicit the desired balance between oscillation and obstacle over-avoidance.

In Experiment 7 (the natural terrain course where each system is tested 10 times), depth based updating has a quicker mean runtime than translation based forgetting—although, not by a significant amount (using the student’s  $t$ -test with the null hypothesis that the two systems’ runtimes are sampled from the same normal distribution gives a  $p$ -value of 0.26). We attribute this (non-significant) performance difference to the slight overestimation of obstacle size by the depth based updating system. If the cylindrical planner is used as the local planning subsystem of a hierarchical planner, then forgetting enabled oscillation can be avoided by using appropriately close subgoals. However, depth updating has no parameters to tune, is less susceptible to translational oscillation, and will never overestimate the size of an obstacle—although its use is limited to systems with depth data.

#### 8.4 Planning systems: image, cylindrical, hierarchical

The wall and V courses (Experiments 2 and 3) show that the basic image planner is susceptible to rotational oscillation whenever an obstacle wider than the FOV obscures the goal. The cylindrical planner is immune to rotational oscillation, hence, it is more robust than the basic image planner. Although the latter can succeed in a subset of environments navigable by the former (i.e. sparsely populated environments containing obstacles narrower than the FOV), it is still advisable to use the cylindrical planner whenever possible. Both the image planner and the cylindrical planner exhibit translational oscillation when a wide obstacle obscures the goal. As discussed in Section 8.2, this is manifested differently depending on the distance metric being used ( $D_i^{\text{S}}$  is less susceptible than  $D_i^{\text{R}^f}$ ). A separate type of translational oscillation, discussed in Section 8.3, is enabled by the translation based forgetting cylindrical updating method.

All three forms of oscillation observed in our experiments are triggered when a large obstacle exists directly between the robot and the goal. This reinforces our belief that image planning is more suited to near-field (local) planning than far-field (global) planning. In order to provide quick navigation commands for obstacle avoidance, any local planner must limit the size of its world view. Therefore, any local planner is vulnerable to oscillation from obstacles larger than its world model. When the cylindrical planner is used as a local subsystem of a hybrid hierarchical planner, the global planner can maximize the effectiveness of the local planner by feeding it sub-goals closer than its minimum oscillation amplitude. This ensures that pathological oscillation will never occur. Further motivation for using the cylindrical planner locally is provided by the fact that, due to perspective, the accuracy of an image based representation of the world is maximized in the near-field.

The natural terrain courses (Experiments 4-7) validate our hypothesis that the hybrid hierarchical planner will outperform the Cartesian hierarchical planner—but only for the specific planners using the  $D_i^{\mathbf{R}^f}$  distance metric. Results for planners using the  $D_i^{\mathbf{S}}$  metric are inconclusive. This illustrates the importance of choosing the correct representation of distance when planning in a non-rectilinear projection of the world. With respect to Experiment 7, using the student’s  $t$ -test with the null hypothesis that two systems’ runtimes are sampled from the same normal distribution, the  $D_i^{\mathbf{R}^f}$  depth based updating hybrid hierarchical planner vs. the Cartesian hierarchical planner gives a  $p$ -value of 0.0017, while the  $D_i^{\mathbf{R}^f}$  translation based forgetting hybrid hierarchical planner vs. the Cartesian hierarchical planner gives a  $p$ -value of 0.050. It is important to note that the hybrid and baseline hierarchical planners use an identical global planner, maintain local maps with similar memory requirements, and achieve comparable local frame-rates. We attribute the positive performance of the  $D_i^{\mathbf{R}^f}$  hybrid hierarchical planner to the local cylindrical planner’s high resolution near-field view. This enables accurate robot localization vs. objects in the FOV, and gives the local cylindrical planner the ability to disambiguate, track, and narrowly avoid obstacles that are distorted by the local Cartesian planner.

## 9 Conclusion

Image space path planning is a planning technique used to discover paths through the image output of a grid-based sensor. We present an image space planning technique for local path planning in unknown unstructured outdoor environments. The basic method involves creating an occupancy grid map directly from image data, then using a graph-search algorithm to find a minimum-cost path through a graph representation of the cost-map. Once found, the graph path is used for navigation in the real world. Our method differentiates between the cost values that are stored in the cost map, the cost values assigned to graph edges, and the distances between cost map elements. Borrowing terms derived from a physical analogy, we label these as force, work, and distance, respectively

We propose three image based planning systems, and perform experimental evaluation with a real robot on eight courses. We find that our method of minimizing work, given explicitly separated notions of force and distance, is robust to various force and distance metrics. Evaluation is performed with two different force metrics ( $F_i^d$  and  $F_i^c$ ) and two different distance metrics ( $D_i^S$  and  $D_i^{R^f}$ ).

The greater usable range of color data, as compared to stereo disparity, gives our naive force from color metric  $F_i^c$  a reactive advantage over our force from disparity metric  $F_i^d$ . Although the former is only used as a proof of concept, it demonstrates that accurate force from color data can improve the performance of an image space planning system by enabling it to detect obstacles sooner.

The distance metric based on the  $L^2$  norm in image space ( $D_i^S$ ) postpones trajectory adjustments into the far field, while the flat world distance metric ( $D_i^{R^f}$ ) prefers trajectory adjustments in the near field. Consequently,  $D_i^{R^f}$  causes the robot to follow straighter paths than  $D_i^S$ , while  $D_i^S$  has a greater oscillation amplitude than  $D_i^{R^f}$  (oscillation amplitude is the distance along the edge of a goal blocking obstacle that the robot will travel before turning around). A relatively large oscillation amplitude gives  $D_i^S$  an advantage in environments

containing obstacles with radii larger than the oscillation amplitude of  $D_i^{\mathbf{R}^f}$  (but smaller than the oscillation amplitude of  $D_i^{\mathbf{S}}$ ).

The basic image planner is susceptible to two types of oscillation: rotational oscillation caused by the goal moving outside the FOV, and translational oscillation caused by the apparent location of the image space goal shifting as a function of robotic movement. The simulated  $2\pi$  radian FOV maintained in the cylindrical planner eliminates rotational oscillation; however, it adds the complication of updating data that has moved outside the camera’s FOV. Despite the cylindrical planner’s immunity to rotational oscillation, it is still susceptible to translational oscillation. Of the two cylindrical updating techniques evaluated, translation based forgetting can be used in conjunction with any type of sensor input. However, it must be tuned to balance an additional form of translational oscillation failure vs. moving unnecessarily far around small obstacles. In contrast, depth based updating does not require tuning, does not enable additional oscillation, and will never overestimate the length of an obstacle—but requires depth data.

The hybrid hierarchical planner avoids all forms of oscillation by giving the cylindrical planner appropriate subgoals from a global Cartesian planner. This framework incorporates the strengths of both planning paradigms, allowing for better overall path planning. Hybrid hierarchical planners using the  $D_i^{\mathbf{R}^f}$  distance metric are able to outperform a double 2D top-down Cartesian planner on four real world courses.

Image space planning provides a convenient framework for maintaining a high-resolution view of near-field terrain by performing path search at the maximum resolution of the image sensor. Assuming an image sensor of fixed resolution, other advantages include: guarantees on search-time and a stable memory footprint. A natural application for image space planning is as the local planning component of a hybrid hierarchical planner. This combination provides the high-resolution near-field view, precise robot vs. obstacle localization, and dependable frame rate of a local image space planner with the translational memory and

fixed-resolution far-field view of a global 2D Cartesian planner. Overall, we find that image space planning is a viable technique for discovering and following paths through unknown unstructured outdoor environments.

## Acknowledgments

This work has been supported by the DARPA LAGR program (DOD AFRL award no. FA8650-07-C-7702) and the National Science Foundation (NSF IIS-0535269). We would like to thank the members of the Intelligence in Action Laboratory at the University of Colorado at Boulder, our fellow participants in the LAGR program, and the LAGR government team. We would also like to thank the anonymous experts that reviewed this work, their comments and suggestions have made it a much better paper.

## References

- Argyros, A., Bekris, K. E., Orphanoudakis, S. C., and Kavraki, L. E. (2005). Robot homing by exploiting panoramic vision. *Autonomous Robots*, 19:7–25.
- Borenstein, J. and Koren, Y. (1991). The vector field histogram—fast obstacle avoidance for mobile robots. *IEEE Journal of Robotics and Automation*, 7:278–288.
- Carsten, J., Rankin, A., Ferguson, D., and Stentz, A. (2007). Global path planning on board the mars exploration rovers. In *IEEE Aerospace Conference*.
- Chen, S. (1990). A spherical model for navigation and spatial reasoning.
- Cowan, N., Weingarten, I., and Koditschek, D. (2002). Visual servoing via navigation functions. *IEEE Transactions on Robotics and Automation*, 18:521–533.
- Dijkstra, E. W. (1959). A note on two problems in connection with graphs. In *Numerical Mathematics*, volume 1, pages 269–271.
- Elfes, A. (1989). Using occupancy grids for mobile robot perception and navigation. In *IEEE Computer*, page 4657.

- Feddema, J. and Mitchell, O. (1989). Vision-guided servoing with feature-based trajectory generation. *IEEE Transactions on Robotics and Automation*, 5:691–700.
- Ferguson, D. and Stentz, A. (2006). Using interpolation to improve path planning: The field d\* algorithm. *Journal of Field Robotics*, 23:79–101.
- Gat, E., Slack, M., Miller, D. P., and Firby, R. J. (1990). Path planning and execution monitoring for a planetary rover. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '90)*, volume 1, pages 20–25, Cincinnati, USA.
- Gaussier, P., C Joulain, S. Z., Blanquet, J. P., and Revel, A. (1997). Visual navigation in an open environment without map. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'97)*, pages 545–550.
- Goldberg, S. B., Maimone, M. W., and Matthies, L. (2002). Stereo vision and rover navigation software for planetary exploration. In *IEEE Aerospace Conference Proceedings*, Big Sky, MT, USA.
- Grudic, G., Mulligan, J., Otte, M., and Bates, A. (2007). Online learning of multiple perceptual models for navigation in unknown terrain. In *International Conference on Field and Service Robotics (FSR'07)*, Chamonix, France.
- Hart, P., Nilsson, N., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. In *Proc. IEEE Transactions On System Science and Cybernetics (SSC-4)*, pages 100–107.
- Herman, M. (1986). Fast, three-dimensional, collision-free motion planning. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '86)*, pages 1056–1063, San Francisco, California.
- Hong, J., Tan, X., Pinette, B., Weiss, R., and Riseman, E. M. (1991). Image-based homing. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '91)*, pages 620–625, New York.

- Hutchinson, S., Hager, G. D., and Corke, P. I. (1996). A tutorial on visual servo control. *IEEE Transactions On Robotics and Automation*, 12:651–670.
- Jochem, T. M., Pomerleau, D. A., and Thorpe, C. E. (1995). Vision-based neural network road and intersection detection and traversal. In *Proc. IEEE Conference Intelligent Robots and Systems*, volume 3, pages 344–349.
- Kelly, A. (1994). Adaptive perception for autonomous vehicles. *Technical Report CMU-RI-TR-94-18*.
- Kelly, A. (2000). Mobile robot localization from large-scale appearance mosaics. *International Journal of Robotics Research*, 19:1104–1125.
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98.
- Koenig, S. and Likhachev, M. (2002). Improved fast replanning for robot navigation in unknown terrain. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '02)*.
- Kolski, S., Ferguson, D., Bellino, M., and Siegwart, R. (2006). Autonomous driving in structured and unstructured environments. In *IEEE Intelligent Vehicles Symposium*, Lausanne, Switzerland and Pittsburgh, USA.
- Koren, Y. and Borenstein, J. (1991). Potential field methods and their inherent limitations for mobile robot navigation. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '91)*.
- Krishnaswamy, G. and Stentz, A. (1995). Resolution independent grid-based path planning. *Technical Report CMU-RI-TR-95-08*.
- Krogh, B. H. and Thorpe, C. E. (1986). Integrated path planning and dynamic steering control for autonomous vehicles. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '86)*, pages 1664–1669, San Francisco, California.

- Kutulakos, K. N., Lumelsky, V. J., and Dyer, C. R. (1993). Vision guided exploration: A step toward general motion planning in three dimensions. In *Proc. IEEE Conference on Robotics and Automation*, volume 1, pages 289–296.
- Mateus, D., Avina, G., and Devy, M. (2005). Robot visual navigation in semi-structured outdoor environments. In *Proc. IEEE International Conference on Robotics and Automation (ICRA'05)*, pages 4691–4696.
- Matsumoto, Y., Sakai, K., Inaba, M., and Inoue, H. (2000). View-based approach to robot navigation. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'00)*, volume 3, pages 1702–1708.
- Matthies, L. (1992). Passive stereo range imaging for semi-autonomous land navigation. *Journal of Robotic Systems*, 9:787–816.
- Minguez, J. and Montano, L. (2004). Nearness diagram (nd) collision avoidance in troublesome scenarios. *IEEE Transactions On Robotics and Automation*, 20:45–59.
- Murray, D. and Jennings, C. (1998). Stereo vision based mapping and navigation for mobile robots. In *Proc. IEEE International Conference on Robotics and Automation (ICRA'97)*, pages 1694–1699, New Mexico.
- Murray, D. and Little, J. (1998). Using real-time stereo vision for mobile robot navigation. In *Proc. of the IEEE Workshop on Perception for Mobile Agents*, Santa Barbara, California.
- Nabbe, B. (2005). Extending the path-planning horizon. *Ph.D. dissertation*.
- Ollis, M., Huang, W. H., Happold, M., and Stancil, B. A. (2008). Image-based path planning for outdoor mobile robots. In *Proc. IEEE International Conference on Robotics and Automation (ICRA'08)*.
- Otte, M. (2007). Path planning in image space for the autonomous navigation of unmanned vehicles in unstructured outdoor environments. *M.S. dissertation, University of Colorado at Boulder*.

- Otte, M., Richardson, S., Mulligan, J., and Grudic, G. (2007). Local path planning in image space for autonomous robot navigation in unstructured environments. In *International Conference on Intelligent Robots and Systems (IROS'07)*, San Diego.
- Pomerleau, D. A. (1989). ALVINN: An autonomous land vehicle in a neural network. *Technical Report CMU-CS-89-107*.
- Procopio, M., Mulligan, J., and Grudic, G. (2007). Long-term learning using multiple models for outdoor autonomous robot navigation. In *International Conference on Intelligent Robots and Systems (IROS'07)*, San Diego.
- Sabe, K., Fukuchi, M., Gutmann, J.-S., Ohashi, T., Kawamoto, K., and Yoshigahara, T. (2004). Obstacle avoidance and path planning for humanoid robots using stereo vision. In *Proc. IEEE International Conference on Robotics and Automation (ICRA'04)*.
- Song, D., Lee, H. N., Yi, J., and Levandowski, A. (2007). Vision-based motion planning for an autonomous motorcycle on ill-structured roads. *Auton Robot*, 23:197–212.
- Sonka, M., Hlavac, V., and Boyle, R. (2008). *Image Processing, Analysis and Machine Vision*. Thompson Learning, Toronto, Ont, third edition.
- Stentz, A. (1995). The focussed D\* algorithm for real-time replanning. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Sugiyama, M., Kawano, Y., Niizuma, M., Takagaki, M., Tomizawa, M., and Degawa, S. (1994). Navigation system for an autonomous vehicle with hierarchical map and planner. In *Proc. of the Intelligent Vehicles '94 Symposium*, page 5055.
- Thorpe, C., Herbert, M. H., Kanade, T., and Shafer, S. A. (1988). Vision and navigation for the carnegie-mellon navlab. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10:362–372.
- Tsugawa, S., Yatabe, T., Hirose, T., and Matsumoto, S. (1979). An automobile with artificial intelligence. In *Proc. Sixth International Joint Conference on Artificial Intelligence*, pages 893–895.

- van den Mark, W., Groen, F., and van den Heuvel, J. C. (2001). Stereo based navigation in unstructured environments. In *IEEE Instrumentation and Measurement Technology Conference*, Budapest, Hungary.
- Vidal, R., Shakernia, O., and Sastry, S. (2003). Formation control of nonholonomic mobile robots omnidirectional visual servoing and motion segmentation. In *Proc. IEEE Conference on Robotics and Automation*, pages 584–589.
- Winters, N., Gaspar, J., Lacey, G., and Santos-Victor, J. (2000). Omni-directional vision for robot navigation. In *IEEE Workshop on Omnidirectional Vision (OMNIVIS'00)*, Hilton Head, South Carolina.
- Zhang, H. and Ostrowski, J. (2002). Visual motion planning for mobile robots. *IEEE Transactions on Robotics and Automation*, 18:199–208.