International Journal of Computer Vision 54(1/2/3), 13–34, 2003

A Dynamical Model of Visually-Guided Steering, Obstacle Avoidance, and Route Selection

BRETT R. FAJEN AND WILLIAM H. WARREN
Department of Cognitive and Linguistic Sciences, Brown University
fajenb@rpi.edu
william_warren@brown.edu

SELIM TEMIZER AND LESLIE PACK KAELBLING
Artificial Intelligence Laboratory, Massachusetts Institute of Technology
temizer@ai.mit.edu
lpk@ai.mit.edu

Received April 25, 2001; Revised February 27, 2002; Accepted April 17, 2002

Abstract. Using a biologically-inspired model, we show how successful route selection through a cluttered environment can emerge from on-line steering dynamics, without explicit path planning. The model is derived from experiments on human walking performed in the Virtual Environment Navigation Lab (VENLab) at Brown. We find that goals and obstacles behave as attractors and repellors of heading, the direction of locomotion, for an observer moving at a constant speed. The influence of a goal on turning rate increases with its angle from the heading and decreases exponentially with its distance; the influence of an obstacle decreases exponentially with angle and distance. Linearly combining goal and obstacle terms allows us to simulate paths through arbitrarily complex scenes, based on information about obstacles in view near the heading direction and a few meters ahead. We simulated the model on a variety of scene configurations and observed generally efficient routes, and verified this behavior on a mobile robot. Discussion focuses on comparisons between dynamical models and other approaches, including potential field models and explicit path planning. Effective route selection can thus be performed on-line, in simple environments as a consequence of elementary behaviors for steering and obstacle avoidance.

Keywords: visual control of locomotion, optic flow, obstacle avoidance, path planning, robot navigation

Humans and other animals have a remarkable ability to coordinate their actions with complex, changing environments. This ability is particularly evident in fundamental behaviors such as prehension and locomotion. With little conscious effort, we routinely reach or walk through cluttered scenes, avoiding obstacles, reaching goals, and intercepting moving targets safely and effectively. The problem of adapting behavior to complex environments has proven a challenge in robotics. Recent trends in behavior-based robotics have taken inspiration from biological solutions to such control problems, particularly those of arthropods, regarding both the architecture of action systems (Brooks, 1986; Pfeiffer et al., 1994; Ritzmann et al., 2000) and principles of sensory control (Duchon et al., 1998; Franceschini et al., 1992; Srinivasan and Venkatesh, 1997). In the present paper, we apply a dynamic model of visually-guided locomotion in humans to the problems of steering, obstacle avoidance, and route selection in mobile robots.

A common approach to robot control is to divide the task into modules that perform peripheral and central subtasks, characterized by Brooks (1991) and Moravec (1981) as "sense, model, plan, act." Under this type of model-based control, sensory information is used to compute a fairly comprehensive internal model of the 3D layout and motions of objects and surfaces in the scene. On the basis of the model, an action path through the scene is then explicitly planned, prior to its execution in the physical environment. The path planning process may compute an optimal route on the basis of some variational principle, such as shortest path, least energy, or minimum jerk. An alternative approach, originating with Gibson (1958/1998, 1979), Lee (1980), and Warren (1988), has sought to achieve on-line control on the basis of occurrent sensory information, without an explicit world model or path planning process (Aloimonos, 1993; Brooks, 1991; Duchon et al., 1998). Under this sort of information-based control, adaptive behavior is governed by mappings between informational variables and action variables, referred to as *laws of control*. One aim of this approach is to determine how apparently planned behaviors such as route selection might emerge as a consequence of the way information is used to modulate action, blurring the line between purely reactive and planned behavior. It remains to be seen how far such an approach can be extended to more complex navigation problems.

A further step in this development has been the introduction of dynamical principles to achieve both stability and flexibility in behavior (Beer, 1995; Schöner and Dose, 1992; Schöner et al., 1995; Warren, 1998b), building upon research in human motor coordination (Kelso, 1995; Kugler and Turvey, 1987). In our version of such an approach, the agent and its environment can be described as a pair of dynamical systems, which are coupled mechanically and informationally. Change in the state of the environment is a function of its current state and forces exerted by the agent, according to the laws of physics; reciprocally, change in the state of the agent (action) is a function of its current state and information about the environment, according to laws of control. Behavior arises from interactions between the components of this mutually coupled system and reflects the constraints of both components. Such systems can be formally described in terms of a set of differential equations, with observed behavior corresponding to solutions to the equations for a given set of initial conditions. Stable modes of behavior and flexible transitions between them are expressed in the low-dimensional dynamics of the system, which we term the *behavioral dynamics*.[1]

More specifically, the emergent behavior can be characterized in terms of *behavioral variables* that are selected on the basis of their relevance to the task goals (Schöner et al., 1995). The current state of the system, as well as intended and avoided states, are thus expressible as (sets of) points in the space of behavioral variables, and behavior corresponds to trajectories through this space. Expressing the behavior in terms of a system of differential equations allows us to exploit the tools and concepts of dynamical systems theory (Strogatz, 1994). In the language of dynamical systems, points toward which trajectories converge are called *attractors* and points away from which trajectories diverge are called *repellors*. Sudden changes in the number or type of such fixed points as a consequence of continuous changes in system parameters can be described as *bifurcations*. Thus, attractors and repellors in the space of behavioral variables may correspond to goal states and avoided states, and bifurcations to qualitative transitions between behavioral modes, providing flexibility in behavior.

We thus distinguish two levels of analysis: the agent-environment interaction (information and control laws), and the emergent behavior (behavioral dynamics). Given that behavior is a complex product of the mutually coupled system, it cannot be simply dictated by the agent. The challenge for the agent or engineer thus becomes one of identifying control laws that evoke the desired behavior in the system as a whole, such that available information is used to shape the appropriate behavioral dynamics.

In this paper we investigate visually-guided locomotion in such a dynamical framework, inspired by the work of Schöner et al. (1995). Our approach will be to identify a set of behavioral variables for steering and obstacle avoidance, measure human behavior when walking to a goal and around an obstacle, and develop a model of the behavioral dynamics. Our ultimate aim is to determine whether successful route selection, in which an agent must detour around one or more obstacles to reach a goal, can be accounted for by the on-line steering dynamics, without an explicit world model or path planning. We then apply the model to the problem of robotic control in simulation and on a mobile robot, and show that it compares favorably with the potential field method of path planning. The model is relevant to route selection through relatively simple

scenes in which the locations of goals and obstacles are currently accessible. Navigation in complex environments (e.g. mazes) is likely to require more sophisticated strategies based on more global knowledge of the environment.

A Dynamical Framework for Steering and Obstacle Avoidance

Consider an agent moving through a simple environment with a constant speed s and a direction of locomotion ϕ, which we will refer to as the *heading*, defined with respect to a fixed allocentric reference axis (Fig. 1). From the agent's current position, a goal lies in the direction ψ_g at a distance d_g, and an obstacle lies in the direction ψ_o at a distance d_o. To steer toward the goal, the agent must turn its heading in the direction of the goal, such that $\phi = \psi_g$ and $\dot{\phi} = 0$. At the same time, the agent must turn away from the obstacle, such that $\phi \neq \psi_o$ when $\dot{\phi} = 0$. Thus, the intended state of steering toward the goal can be expressed by particular values of ϕ and $\dot{\phi}$, and the avoided state of steering toward the obstacle can be expressed by different values of ϕ and $\dot{\phi}$. Because ϕ and $\dot{\phi}$ provide a set of variables that can be used to express the current state of the system, as well as intended and avoided states, we adopt ϕ and $\dot{\phi}$ as behavioral variables.

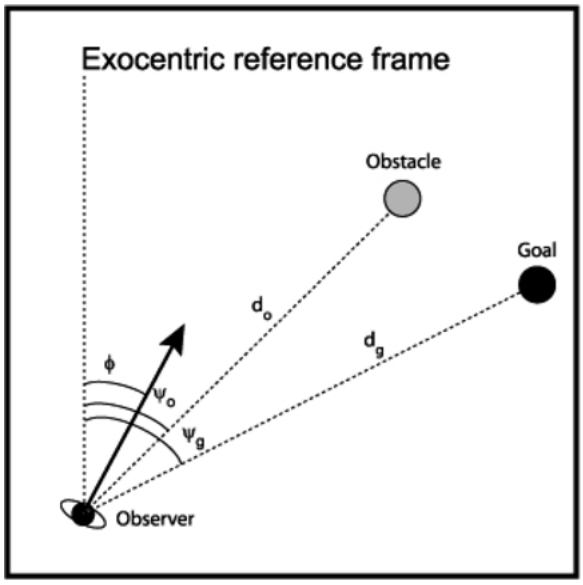

Figure 1. Plan view of an observer moving through an environment containing a goal and an obstacle. The dotted line is a fixed, exocentric reference line used to define the observer's direction of locomotion (ϕ), the direction of the goal (ψ_g) and the direction of the obstacle (ψ_o). d_g and d_o correspond to the distance from the observer to the goal and obstacle, respectively.

We next develop a model in the form of a system of differential equations that describes how the behavioral variables change over time, analogous to a mass-spring system. Broadly speaking, the model consists of three components: a goal component, an obstacle component, and a damping term. The damping term opposes turning, and we assume it is a monotonically increasing function of $\dot{\phi}$ and is independent of ϕ. The goal component determines how the egocentric location of a goal contributes to angular acceleration ($\ddot{\phi}$), and is assumed to be a function of the current *goal angle* ($\phi - \psi_g$) and goal distance (d_g). Finally, the obstacle component determines the contribution of each obstacle in the scene and is assumed to be a function of the *obstacle angle* ($\phi - \psi_o$) and possibly obstacle distance (d_o). Taken together, the general form of the model is:

$$\ddot{\phi} = -f_d(\dot{\phi}) - f_g(\phi - \psi_g, d_g) + \sum_{i=1}^{\#obstacles} f_o\left(\phi - \psi_{o_i}, d_{o_i}\right) \quad (1)$$

where f_d is the damping function, f_g is the goal function, f_o is obstacle function, and the subscript i is the index of each obstacle in the scene. Although the motion of the agent to a new (x, z) position in the environment will alter ψ_g, d_g, ψ_o, and d_o, these variables can be rewritten as functions of x and z (see Appendix). The agent-environment system is thus completely described by a four-dimensional system of equations, for to predict the agent's future position we need to know its current position (x, z), heading (ϕ), and turning rate ($\dot{\phi}$), assuming that speed is constant. (See Appendix for the complete set of equations.) Note, however, that at this stage the agent and objects are simply treated as points.

The precise manner in which the agent turns toward goals and away from obstacles is determined by the form of each function, and reflected in the shape of the trajectory through the space of behavioral variables. To select the form of each function, we turned to empirical observations of human walking. We designed a series of experiments intended to measure how the angles and distances to goals and obstacles influence the turning rate. These observations were then used to specify the form of goal and obstacle functions and estimate parameter values in the dynamical model.

Human Experiments

Three experiments were designed to reveal the factors that influence how humans turn toward goals and away from obstacles during walking (see Fajen and Warren, 2003), for details). The studies were conducted in the Virtual Environment Navigation Lab (VENLab) at Brown University. The VENLab consists of a 12 m × 12 m room in which subjects are able to walk around freely while wearing a head-mounted display (HMD). A hybrid inertial and ultrasonic tracker mounted in the ceiling tracks the position and orientation of the HMD. This information is fed back to a high-performance graphics workstation, which updates the visual display presented in the HMD. This facility allows us to manipulate both the structure of the environment and the visual information presented to the observer in real-time, while simultaneously recording ongoing behavior in naturalistic tasks.

The first experiment examined the simple case of walking toward a goal, while the second examined avoiding a single obstacle en route to a goal. In Experiment 1, observers began each trial by walking in a specified direction. After walking 1 m, a goal appeared at an angle of $\phi - \psi_g = 5^\circ$, 10°, 15°, 20°, or 25° from the heading direction and a distance of $d_g = 2, 4$, or 8 m. Observers were simply asked to walk to the goal. The major findings of Experiment 1 were that the turning rate and angular acceleration toward goals increased with goal angle (see Fig. 2(a)) but decreased with goal distance (see Fig. 2(b)). In Experiment 2, observers began walking toward a goal located straight ahead at a distance of 10 m. After walking 1 m, the obstacle appeared at an angle of $\phi - \psi_o = 1^\circ, 2^\circ, 4^\circ$, or 8° from the heading direction and a distance of $d_o = 3, 4$, or 5 m. The major findings of Experiment 2 were that the turning rate and angular acceleration away from obstacles decreased with both obstacle angle (see Fig. 3(a)) and obstacle distance (see Fig. 3(b)).

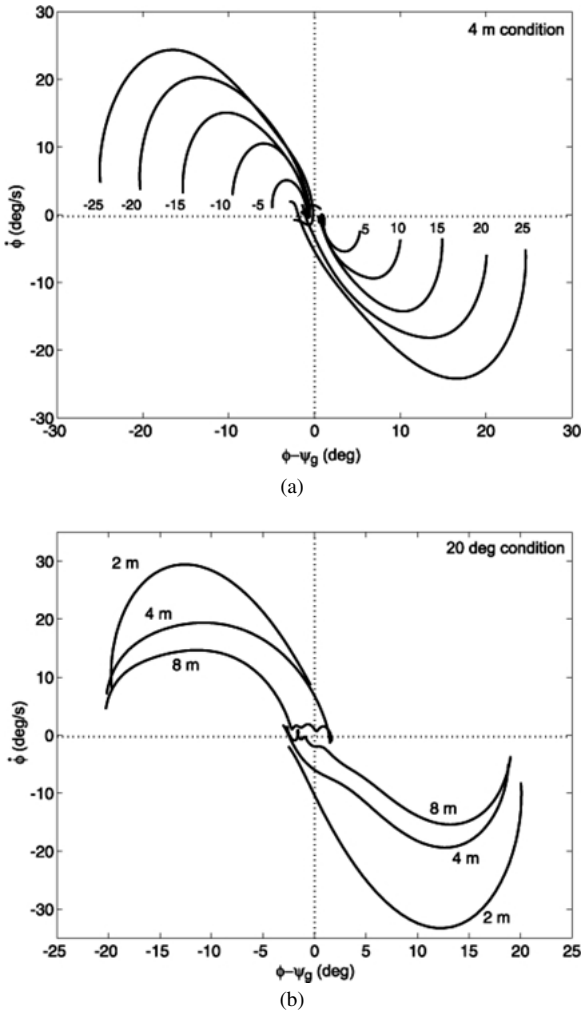

Figure 2. Human trajectories for turning toward a goal in Experiment 1 (turning rate ($\dot{\phi}$) vs. goal angle ($\phi - \psi_g$)). Curves correspond to (a) different initial goal angles in the 4 m condition and (b) different initial goal distances in the 20° condition.

The Model

These empirical observations were used to specify the dynamical model of steering and obstacle avoidance. First, for purposes of simplicity, we assumed that damping would be proportional to turning rate, such that $f_d(\dot{\phi}) = b\dot{\phi}$, for some constant $b > 0$. The goal function $f_g(\phi - \psi_g, d_g)$ was chosen to reflect the findings that the influence of the goal on angular acceleration increases with goal angle and decreases with goal distance:

$$f_g(\phi - \psi_g, d_g) = k_g(\phi - \psi_g)(e^{-c_1 d_g} + c_2) \qquad (2)$$

Thus, in the model the goal's influence increases *linearly* with goal angle up to 180° (see Fig. 4(a)) and decreases *exponentially* with goal distance (see Fig. 4(b)). Note that this influence asymptotes to some minimum non-zero value as goal distance increases, enabling the agent to steer toward distant goals. The "stiffness" parameter k_g is a gain term for the goal component, c_1 sets the rate of exponential decay with goal distance, and c_2 scales the minimum acceleration toward distant goals.

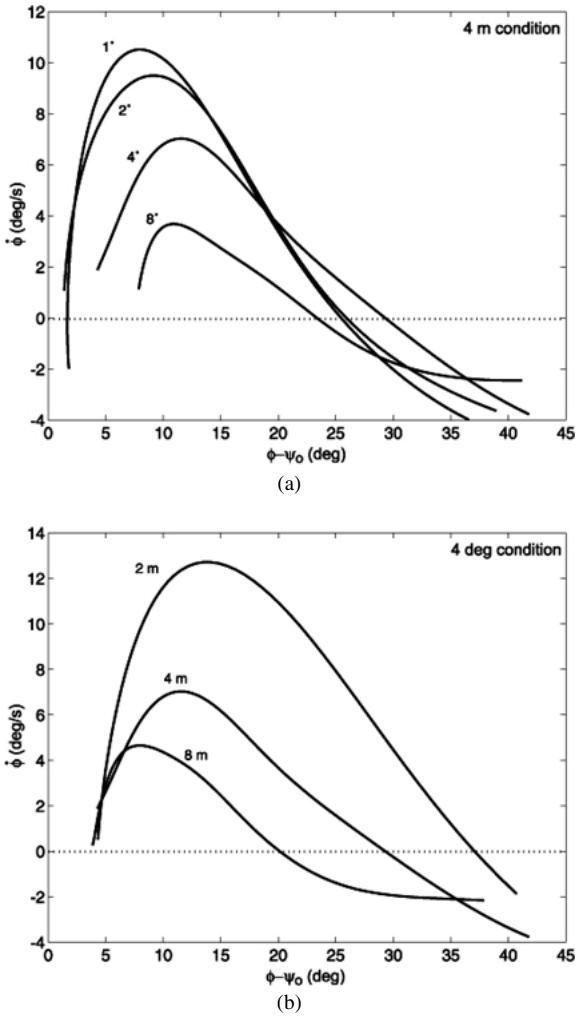

Figure 3. Human trajectories for turning away from an obstacle in Experiment 2 (turning rate ($\dot{\phi}$) vs. goal angle ($\phi - \psi_g$)). Curves correspond to (a) different initial obstacle angles in the 4 m condition and (b) different initial obstacle distances in the 4° condition.

Likewise, the obstacle function $f_o(\phi - \psi_o, d_o)$ was chosen to reflect the findings that the influence of the obstacle on angular acceleration decreases with both obstacle angle and distance:

$$f_o(\phi - \psi_o, d_o) = k_o(\phi - \psi_o)\left(e^{-c_3|\phi-\psi_o|}\right)(e^{-c_4 d_o}) \quad (3)$$

In this case, the obstacle's influence decreases *exponentially* with obstacle angle (see Fig. 4(c)) as well as with obstacle distance (see Fig. 4(d)). The parameter k_o is a gain term for the obstacle component, c_3 sets the rate of decay with obstacle angle, and c_4 sets the rate of decay with obstacle distance. Note that for small obstacle angles, acceleration away from the obstacle increases with obstacle angle, such that the function is continuous and there is a repellor at an obstacle angle of zero. Unlike the goal component, the obstacle influence decreases to zero as distance goes to infinity. When parameterized to fit the human data, these two exponentials imply that only obstacles within ±30° of the heading direction and less than 4 m ahead exert an appreciable influence on steering behavior. Note that the exponential terms introduce nonlinearity into the system.

Thus, the full model is:

$$\ddot{\phi} = -b\dot{\phi} - k_g(\phi - \psi_g)(e^{-c_1 d_g} + c_2) + k_o(\phi - \psi_o)\left(e^{-c_3|\phi-\psi_o|}\right)(e^{-c_4 d_o}) \quad (4)$$

In principle, additional obstacles in the environment can be included by simply adding terms to the equation. The model thus scales linearly with the complexity of the scene, and doesn't blow up in complicated environments (Large et al., 1999). Furthermore, only obstacles near the heading direction and a few meters ahead need to be evaluated, making the model computationally quite tractable. The agent therefore does not need a memory representation of the entire scene; as long as the goal location is available to the agent's sensors, route selection is performed simply on the basis of the obstacles within a small spatial window ahead.

Simulations

We simulated the model under a variety of conditions to test its success in steering toward goals, avoiding obstacles and selecting routes. The conditions used for the first two sets of simulations were identical to those used in the two preceding human experiments, and their purpose was to test the adequacy of Eq. (4) as a model of human behavior. The next step was to test the model in more complex scenes containing one or more obstacles in which multiple routes around the obstacle(s) are possible. These simulations were intended to reveal how goal and obstacle components interact to perform route selection.

Simulation #1: Steering Toward a Goal

We simulated the model under the same conditions used in Experiment 1 on steering toward a goal, to identify the single set of parameters for the goal component that best fit the data. Simulations were compared with the mean time series of goal angle in the human data

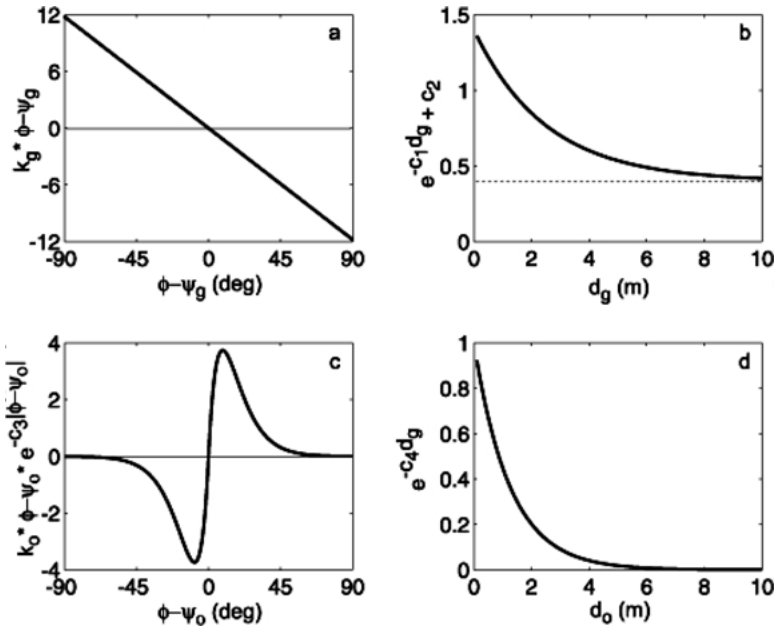

Figure 4. Plots of (a) goal angle term, (b) goal distance term, (c) obstacle angle term, and (d) obstacle distance term from Eq. (4).

using a least-squares analysis, as the four parameters were systematically varied. The best fit ($r^2 = 0.982$) was found with parameter values of $b = 3.25$, $k_g = 7.50$, $c_1 = 0.40$, and $c_2 = 0.40$. Using these settings, the model produced paths to the goal that were virtually identical with human subjects (Fig. 5), turning at a rate that depended on goal angle and distance in a similar manner. Specifically, turning rate and angular acceleration increased with goal angle (Fig. 6(a)) and decreased with goal distance (Fig. 6(b)).

Simulation #2: Avoiding an Obstacle

Adding a single obstacle component, we simulated the model under the conditions used in Experiment 2. We used the parameter settings found in the previous simulation for the goal component, and fit the three parameters for the obstacle component in the same manner as before. The best fitting obstacle values (mean $r^2 = 0.975$) were $k_o = 198.0$, $c_3 = 6.5$, and $c_4 = 0.8$. Using these settings, the model successfully detoured around the obstacle to the goal on paths very similar to those of human subjects (Fig. 7). The turning rate and acceleration away from the obstacle decreased with obstacle angle (see Fig. 8(a)) and decreased with obstacle distance (see Fig. 8(b)), reproducing the characteristics of human obstacle avoidance behavior. Thus, the model exhibits both a good quantitative and qualitative fit to the human behavior observed in Experiments 1 and 2.

Simulation #3: Route Selection

To see whether the model could predict the routes humans would select through somewhat more complex scenes, we performed simulations with a variety of other goal and obstacle configurations. Because the model functions in real-time, behavior is determined entirely by the interaction of goal and obstacle components, whose influence changes with the position, heading and turning rate of the agent. How might goal and obstacle components interact to determine the route?

Simulation #3a: Relative Position of Goal and One Obstacle. Consider the situation in which the direction of the obstacle lies in between the direction of heading and the direction of the goal (see Fig. 9). In this case, the agent could take either an outside (left) path or an inside (right) path around an obstacle. If the agent's behavior is determined by the interaction of goal and obstacle components, and if the relative "attraction" of the goal and "repulsion" of the obstacle depend on their locations, then the offset angle between the obstacle and goal and the goal distance should influence the agent's route.

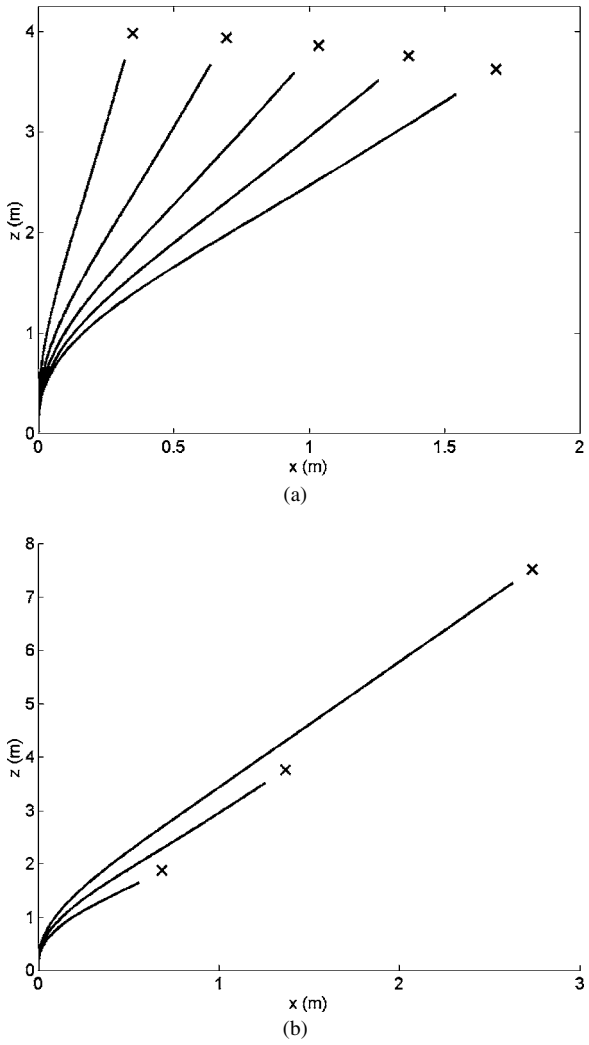

Figure 5. Paths produced by model to goals located at (a) 5°, 10°, 15°, 20°, and 25° and 4 m and (b) 2, 4, and 8 m in the 20° condition in Simulation #1.

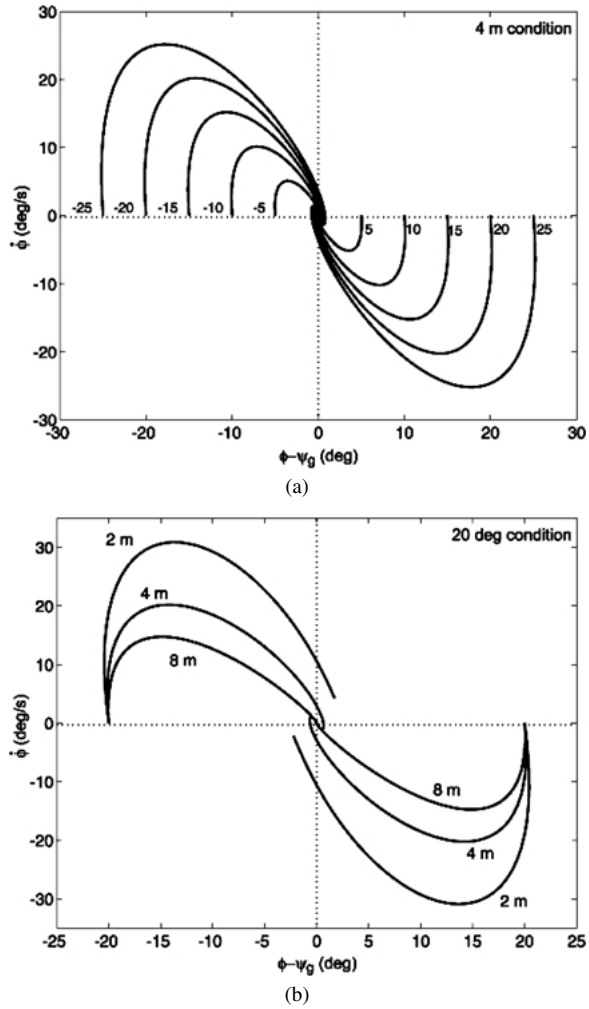

Figure 6. Model trajectories in Simulation #1 (turning rate ($\dot{\phi}$) vs. goal angle ($\phi - \psi_g$)). Curves correspond to (a) initial goal angle in the 4 m condition and (b) initial goal distance in the 20° condition.

We tested the model using configurations of goals and obstacles similar to those in Fig. 9. Keeping the initial goal angle constant at 15° and the initial obstacle distance constant at 4 m, we varied the initial goal distance between 5 m and 9 m, and the initial offset angle between 1° and 15°. We found effects of both initial goal distance and initial offset angle. Using the fixed parameters determined in Simulations #1 and #2, the agent selects an outside route for offset angles $\leq$7°, and an inside path for angles $\geq$10°. For angles between 7° and 10°, the agent takes an outside route for larger goal distances and switches to an inside route for smaller goal distances (Fig. 10).

The effect of initial goal distance is a consequence of the fact that the attractive strength of the goal, and hence angular acceleration toward the goal, increases as the goal gets nearer. The effect of offset angle is a consequence of the trade-off between the attractive strength of the goal, which increases with angle, and the repulsive strength of the obstacle, which decreases with angle. Initially, the goal component dominates, turning the agent in the direction of the goal. The resulting decrease in both goal and obstacle angle decreases the attractive strength of the goal and increases the repulsive strength of the obstacle. Whether the agent follows an inside or outside route depends on which component dominates as the agent heads toward the obstacle. For large offset angles, the goal angle is relatively large

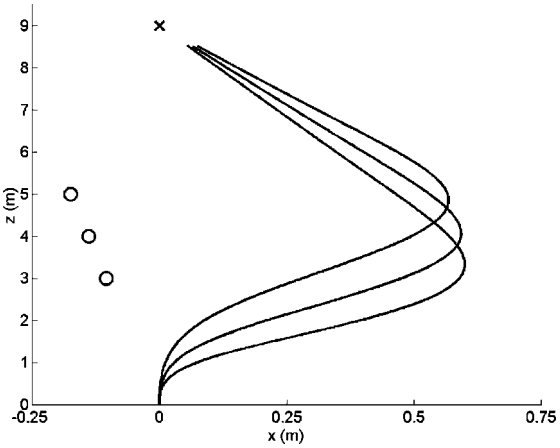

Figure 7. Paths produced by model around obstacles located at 4° and 3, 4 or 5 m in Simulation #2.

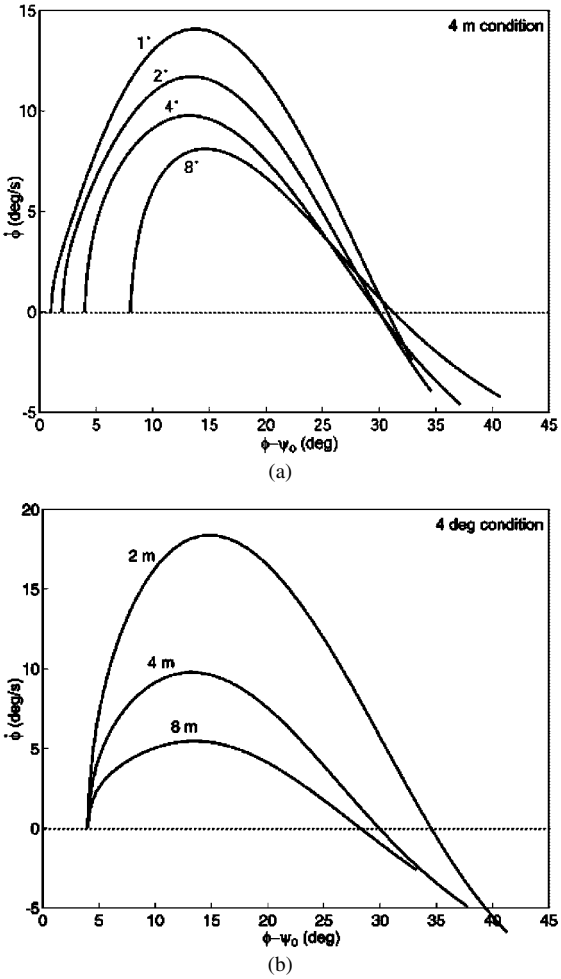

Figure 8. Model trajectories in Simulation #2 (turning rate ($\dot{\phi}$) vs. goal angle ($\phi-\psi_g$)). Curves correspond to (a) initial obstacle angle in the 4 m condition and (b) initial obstacle distance in the 4° condition.

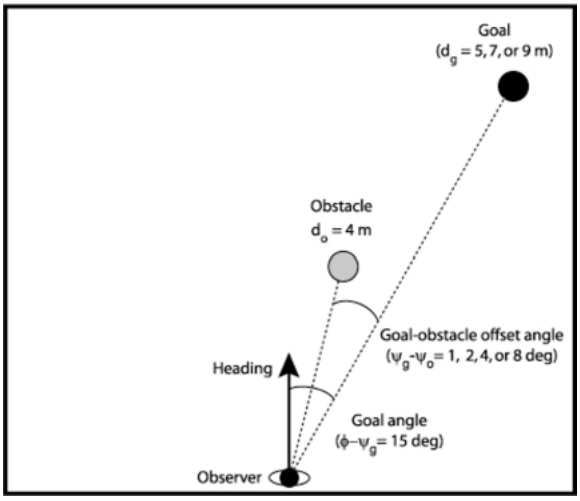

Figure 9. Configuration of goal and obstacle used in Simulation #3a.

as the agent turns toward the obstacle. Hence, goal attraction overcomes obstacle repulsion resulting in an inside route. For small offset angles, the goal angle is relatively small as the agent turns toward the obstacle. Hence, obstacle repulsion overcomes goal attraction, forcing the agent along an outside route. Thus, the deep structure of the observed route selection is represented in the behavioral dynamics.

To evaluate the model's predictive ability, we tested for these effects of initial offset angle and initial goal distance in humans. As in Experiments 1 and 2, subjects began walking in a specified direction. After walking

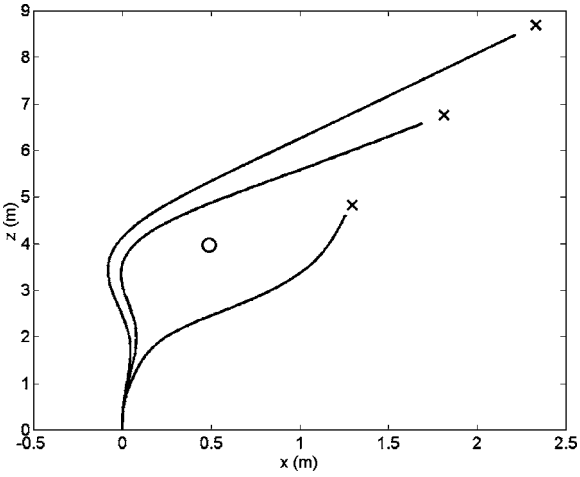

Figure 10. Paths produced by the model to goals located at 15° and 5, 7, or 9 m. Goal-obstacle offset angle is 8° and obstacle distance is 4 m.

1 m, a goal and obstacle appeared simultaneously. Initial goal angle was fixed at 15° and initial obstacle distance at 4 m. We varied initial goal distance between 5, 7, and 9 m and initial offset angle between 1°, 2°, 4°, and 8° (Fig. 9). Means paths for each condition of initial offset angle are shown in Fig. 11(a). Although observers took both inside and outside paths in each condition, the percentage of inside paths decreased with initial goal distance and increased with offset angle (see Fig. 11(b)). Both effects are consistent with the predictions of the model. The distribution of paths could presumably be reproduced by adding a noise term to the model.

Interestingly, the shift to inside paths occurred at somewhat larger offset angles for the model (7–10°) than for human participants (2–4°). Thus, the parameter settings derived from Experiments 1 and 2 yield behavior that is somewhat biased toward outside paths. One reason for this may be that the first two experiments sampled a limited range of conditions, and in particular did not include cases in which participants crossed in front of the obstacle to reach the goal. It is possible that they adapted their behavior (adjusted their "parameters") to these special conditions, with the result that the parameter fits did not generalize precisely to a wider range of conditions. We thus performed a second set of simulations to determine whether we could reproduce the pattern of routes observed in Experiment 3 with a minimal change in parameters. Adjusting a single parameter, c_4, from 0.8 to 1.6, was sufficient to induce the shift from an outside to an inside path at offset angles between 1° and 4°. The c_4 parameter determines the decay rate of obstacle repulsion as a function of distance, and increasing it results in somewhat "riskier" behavior. Thus, the model successfully predicted the qualitative effects of initial goal distance and initial offset angle on route selection, and with a minor adjustment to one parameter reproduced the quantitative properties of the human data.

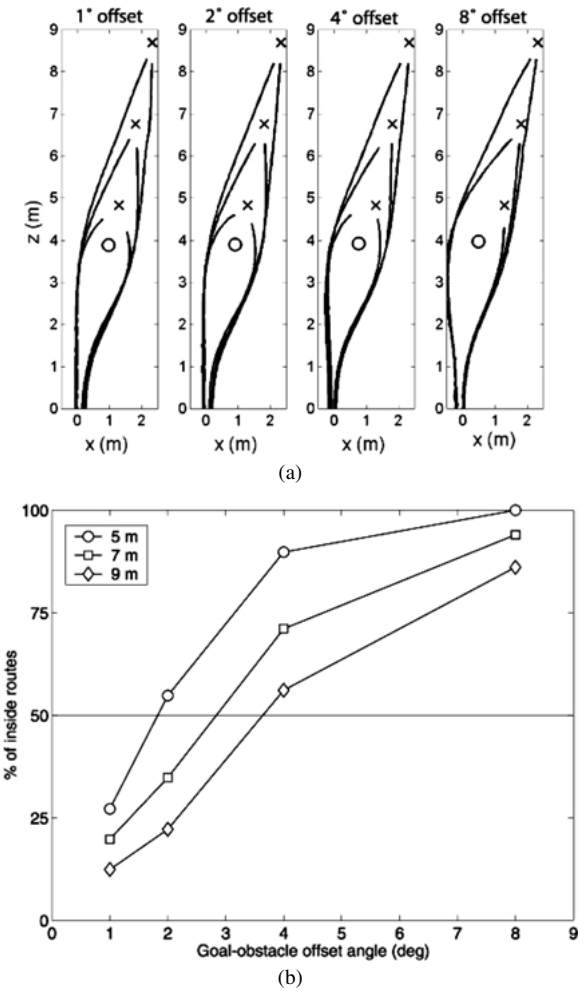

Figure 11. (a) Mean paths and (b) percentage of inside paths produced by humans under conditions used in Simulation #3a.

Simulation #3b: Relative Position of Two Obstacles. Whereas Simulation #3a was intended to reveal how goal and obstacle components interact, Simulation #3b focused on the interaction of two obstacle components. Specifically, we wanted to determine how the location of a distant obstacle affects the agent's route around a nearby obstacle. In this set of simulations, the initial angle (0°) and distance (9 m) of the goal was fixed, as was the initial angle (0.5°) and distance (4 m) of the nearby obstacle. We manipulated the initial angle of the distant obstacle while keeping its initial distance fixed at 4.5 m (Fig. 12(a)). When the angle of the distant obstacle was close to zero (−0.5°), the agent detoured to the left of both obstacles (Fig. 12(b)). As that angle grew slightly (−5°), the agent detoured to the right of both obstacles (Fig. 12(c)). Finally, as the angle opened further (−15°), the agent switched to a route between the two obstacles (Fig. 12(d)).

The agent appears to be making intelligent route selection decisions, choosing the route that is most efficient for the given configuration of obstacles. It is easy to see, however, how these "choices" emerge from the interaction of the two obstacle components. Because the two obstacles are initially on opposite sides of the agent's heading, they oppose one another. When the

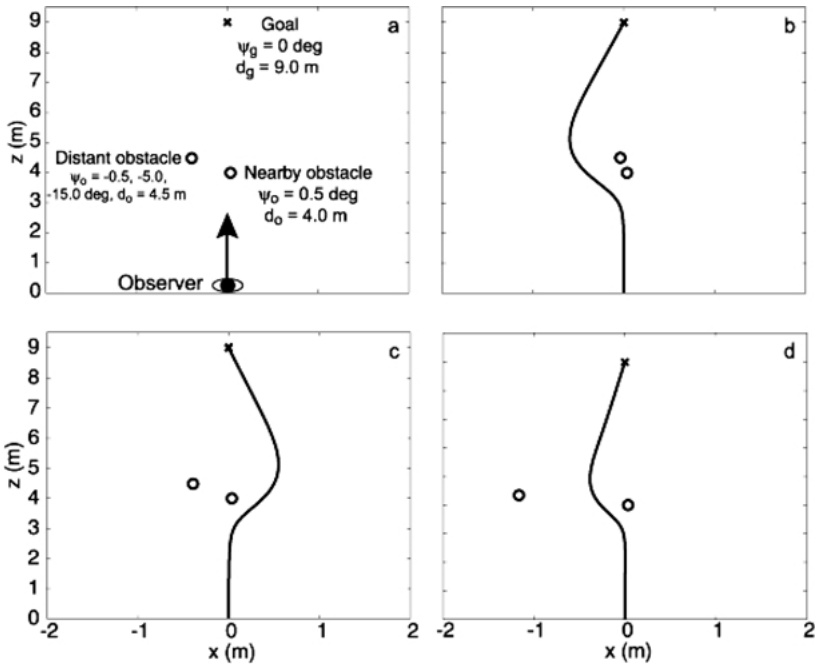

Figure 12. (a) Configuration of goal and obstacles used in Simulation #3b, and (b)–(d) paths produced by the model for distant obstacle angles of -0.5°, -5.0° and -15°, respectively.

initial angle of the distant obstacle is -0.5° (the same as the nearby obstacle, but opposite sign), the nearby obstacle dominates because it is closer. Opening that angle to -5° increases the influence of the distant obstacle because its exponential function grows until about -8° (Fig. 4(c)). Hence, the distant obstacle dominates, forcing the agent to the right of both obstacles. When the angle is changed to -15°, the influence of the distant obstacle decreases again because the exponential function returns toward zero at larger obstacle angles. Early in the run, the nearby obstacle dominates, forcing the agent to the left. This closes the distant obstacle angle, however, increasing its influence until it equals that of the nearby obstacle. Thus, unlike the situation in Fig. 12(b), the agent follows a route in between the two obstacles. In each of these cases, the possible routes at any moment appear as point attractors in the behavioral dynamics, and switching between routes correspond to bifurcations between attractor layouts.

Simulation #3c: Route Selection Through a Field of Randomly Positioned Obstacles. Our ultimate aim was to model route selection through a field of randomly positioned obstacles to reach a goal. Initial goal angle and distance were fixed at 0° and 9 m, respectively. Ten obstacles were then randomly positioned in a rectangular area 4 m wide and 7 m long centered between the agent's initial position and the position of the goal (see Fig. 13(a)). We simulated the model many times, using different random configurations of obstacle on each trial, and found that the agent always reached the goal without colliding into any obstacles. Furthermore, the agent generally followed smooth, efficient routes to the goal, never crossing its own path or getting trapped. Several examples are shown in Fig. 13(b)–(d).

Comparison of the Dynamical Model with Potential Field Methods

Since Khatib's (1986) influential paper, a dominant technique for local obstacle avoidance in mobile robots has been the *potential field method*. In this section, we briefly introduce potential field methods, compare simulations of dynamical and potential field models in a sample environment, and discuss the major differences between the dynamical and potential field methods in some detail.

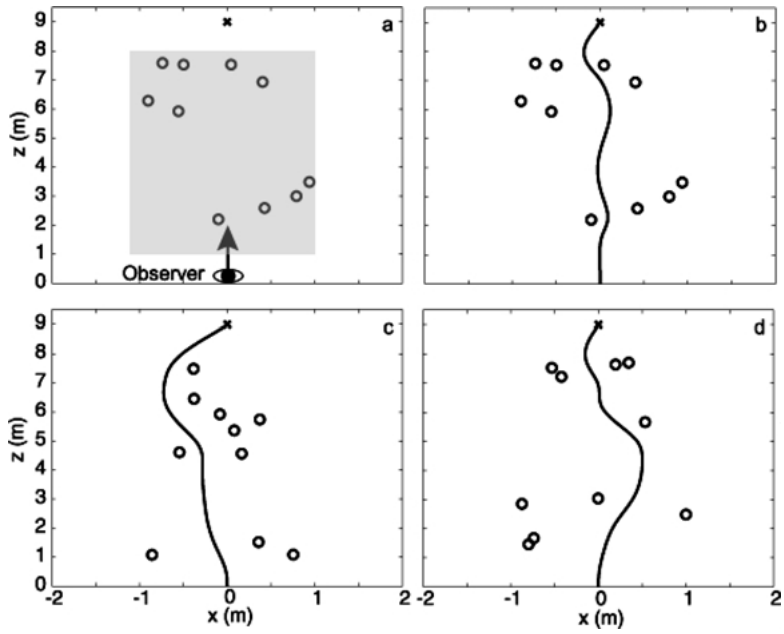

Figure 13. (a) Sample configuration of a field of obstacles used in Simulation #3c. Gray rectangle defines region in which obstacles were randomly placed. (b)–(d) Sample paths produced by the model through a field of obstacles.

Potential Field Methods

The basic concept of potential field methods is the following: Given the task of traveling from a starting location to a target location without running into obstacles, two kinds of imaginary forces are generated that act on the agent's current location $\vec{\chi}$. The first are attractive forces due to an attractive potential field $U_{x_d}(\vec{\chi})$. They originate at the target location and *pull* the agent in the direction of the target. The second are repulsive forces, which are due to a repulsive potential field $U_O(\vec{\chi})$ and originate at obstacles to *push* the agent away. The attractive and repulsive fields are linearly combined in an artificial potential field $U_{\text{art}}(\vec{\chi})$. Using the gradient of this field, we can compute a resultant force vector that is used to control the agent's direction and often speed of motion:

$$U_{\text{art}}(\vec{\chi}) = \underbrace{U_{x_d}(\vec{\chi})}_{\text{Attractive Potential Field}} + \underbrace{U_O(\vec{\chi})}_{\text{Repulsive Potential Field}} \tag{5}$$

In Khatib's (1986) formulation, the attractive potential of the target increases with the square of its distance from the agent,

$$U_{x_d}(\vec{\chi}) = \frac{1}{2} k_p (\vec{\chi} - \vec{\chi}_d)^2 \tag{6}$$

where $\vec{\chi}_d$ describes the target position and k_p the position gain. Conversely, the repulsive potential of an obstacle obeys the inverse square law of distance ρ from the agent, once the agent enters the obstacle's radius of influence ρ_o,

$$U_O(\vec{\chi}) = \begin{cases} \frac{1}{2}\eta\left(\frac{1}{\rho} - \frac{1}{\rho_0}\right)^2, & \text{if } \rho \le \rho_0 \\ 0, & \text{if } \rho > \rho_0 \end{cases} \tag{7}$$

where η is a constant gain.

This basic concept can have slight variations in implementation. For example, both the negative gradients of the above potential fields (Khatib, 1986) and other simpler linear functions (Arkin, 1989) of distance have been used in the computation of force magnitude, and potential functions have been constructed by combining individual obstacle functions with logical operations (Newman and Hogan, 1987). There have also been some enhancements such as taking into consideration the agent's velocity in the vicinity of obstacles (Krogh,

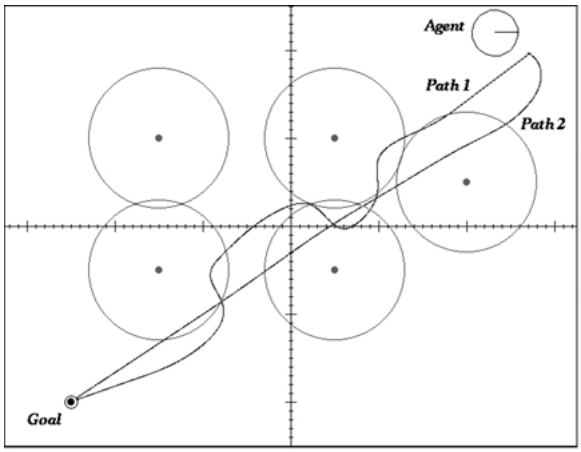

Figure 14. A typical performance example. Large tick marks indicate 1 m intervals.

1984). Potential field methods have been applied to offline path planning (Thorpe, 1985) and in mobile robots with real sensory data (for example by Arkin, 1989).

A Typical Performance Example

We tested both methods in a sample environment containing five obstacles (see Fig. 14), using Khatib's (1986) original potential field formulation. The environment consisted of a 5 m × 6.5 m room with a starting location (indicated by the circle), a target location (labeled goal), and five randomly positioned obstacles (shown as dots). The circles around the obstacles indicate the limit distance of repulsive influence for the potential field model (0.8 m). The agent was assumed to have a diameter of 0.5 m, similar to a human, and an initial heading of 0° (parallel to the x-axis). Although the potential field is often used to control the agent's velocity (direction and speed), in all our simulations we used the resultant force vector to control the agent's direction only, while holding speed constant, analogous to the dynamical model. The straightforward application of the potential field method to mobile robot navigation treats the robot as a particle; however, most mobile robots are non-holonomic, which means they cannot move in arbitrary directions (e.g., without first stopping and turning). In our simulations and robot experiments, we used a controller based on the idea that the front point of a differential-drive robot can be treated as holonomic (Temizer, 2001; Temizer and Kaelbling, 2001). An alternative approach, used by Arkin (1989), for example, is to have the robot repeatedly: stop, turn in the direction of the local force, traverse a short linear segment, stop, reorient, etc. The details of the paths resulting from this method would differ from those we show here, but will be qualitatively similar.

Path 1 shows the trajectory generated by the potential field method, and path 2 (which is almost a straight line) that generated by the dynamical model. In this simulation, the agent moved with a constant translation speed of 0.5 m/s for both methods. Path 1 has a length of 7.55 meters and was traversed in 15.1 seconds, whereas Path 2 was only 6.70 meters long and was traversed in 13.4 seconds. We also implemented the potential field method in a research robot (RWI B21r indoor robot) and we note that the software simulations closely reflect the actual trajectories observed.

The 3D plots in Fig. 15 represent the artificial potential field and the resultant force vectors for the example scene. The top graph (Fig. 15(a)) shows the artificial potential field and the middle graph (Fig. 15(b)) shows the magnitudes of the resultant force vector at each location in the environment, with coordinates that match those of Fig. 14. The starting point is near the high corner, the goal is near the low corner, and the obstacles generate tall cones that extend to infinity, guaranteeing that the agent will never collide with an obstacle.

Differences Between the Two Methods

In this section we consider high-level conceptual differences between the dynamical model and the potential field method. A low-level quantitative comparison would not be appropriate since the computational outcomes of the two methods are quite different: the potential field method produces a resultant vector that directly controls the agent's direction, whereas the dynamical model produces an angular acceleration that controls the agent's rotation.

Angular Acceleration vs. Direction Control. Looking at the example in Fig. 14, it is apparent that the dynamical model tends to traverse smoother and shorter paths than the potential field method. Similarly, the fluctuations in rotation speed are smooth for the dynamical model (Fig. 16), in contrast to sharp, rapid turns with the potential field method. This is partially due to an important general difference between the approaches: the dynamical model explicitly controls the agent's angular acceleration and deceleration rather than the translation direction, and thus tends to generate smoother trajectories. The damping term constrains the

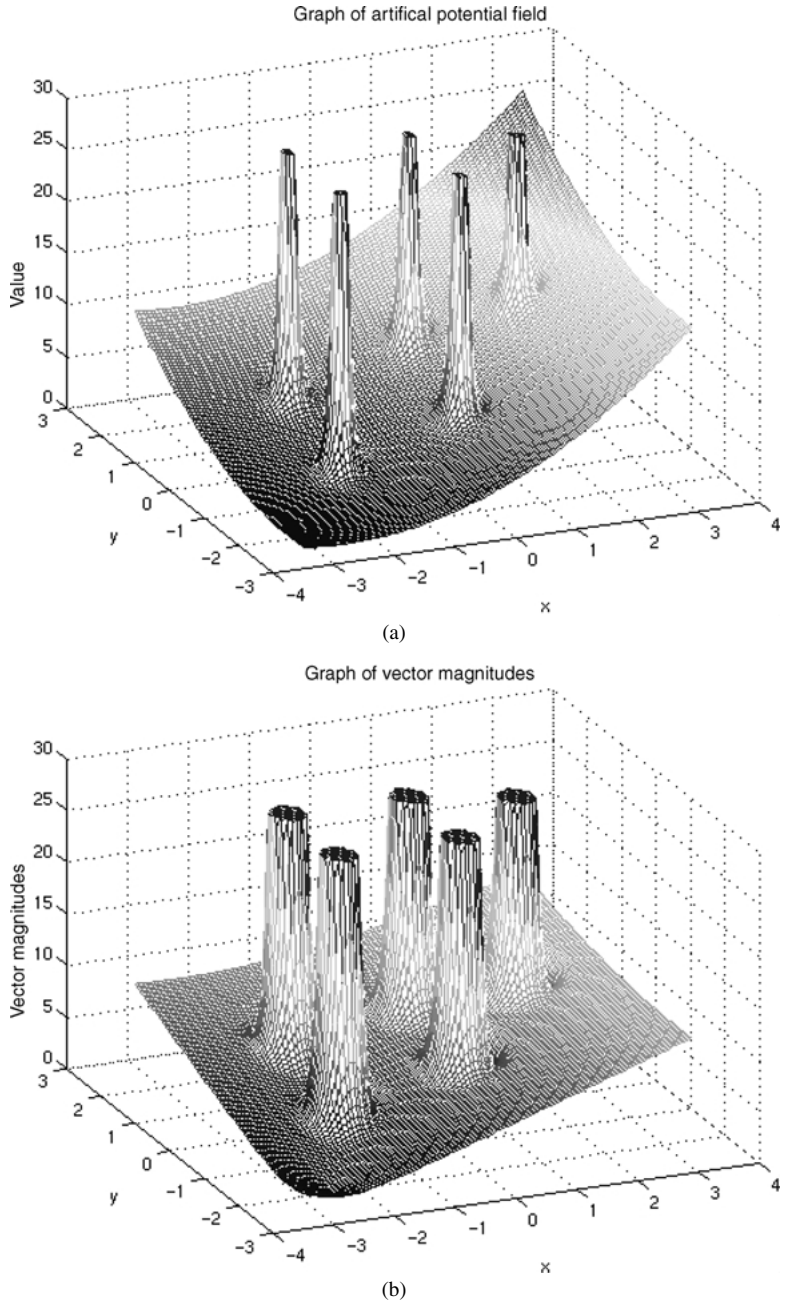

Figure 15. (a) Artificial potential field inside the room and (b) and vector magnitudes.

rotational acceleration, which also acts to smooth the path. In contrast, the potential field method can generate rapid changes in the direction of the velocity vector resulting in frequent sharp turns, depending on the complexity of the artificial potential field (which usually is composed of many hills and valleys even if there are only three or four obstacles; see Fig. 15).

The Obstacle Function. A second reason for smoother, shorter paths stems from another important difference between the two methods. Whereas the effect of the target is similar in both, serving to draw the agent toward the goal, the effect of an obstacle is very different. In the potential field method, the obstacle function depends only on the shortest distance

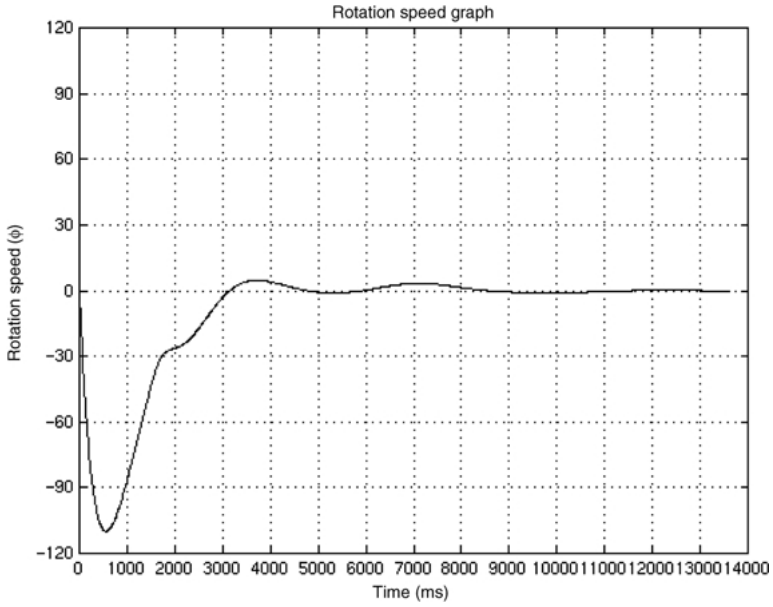

Figure 16. Rotation speed graph for the dynamical model.

between the boundaries of the agent and the obstacle, independent of the agent's heading. In the dynamical model, on the other hand, the obstacle function explicitly depends on the agent's heading relative to the obstacle, as well as obstacle distance. This is apparent in Fig. 17, which plots the obstacle's repulsive influence as a function of its heading angle and distance for each method.

Under the potential field method, once an agent enters the obstacle's circular region of influence, traveling toward the target on a chord of the circle, the magnitude of the repulsive force begins to increase. We can decompose this repulsive force into two orthogonal components: a longitudinal component parallel to the direction of the goal's attractive force, and a lateral component perpendicular to it. As the agent travels along the first half of the chord, the magnitude of the longitudinal component increases more rapidly than the lateral component. But as it approaches the mid-point, the longitudinal component becomes large relative to the attractive force of the goal, and the lateral component increases rapidly, leading the agent to turn away from the obstacle under the influence of the lateral component. For fat (as opposed to point) agents, as the agent passes the obstacle, the lateral component continues to increase and push it farther away, even after the agent has made a sufficient positional correction to avoid the obstacle. This results in typical swerving trajectories like those in Fig. 14.

In contrast, the dynamical model makes effective use of the heading angle with respect to the obstacle ($\phi - \psi_o$), so that the repulsive influence decreases rapidly with obstacle angle as well as distance (Fig. 17(a)). Consequently, once the agent makes a sufficient heading correction, the rotation ceases and it simply passes the obstacle on a smooth path. We also observed that the obstacle angle and obstacle distance terms should be adjusted together, for it is the combined effect of these two parts that produce the smooth and human-like trajectories.

We performed two simulation experiments that demonstrate these effects by manipulating the effective distance of obstacle repulsion. In the first experiment (Fig. 18), the well-balanced parameters of the dynamical model were left intact ($c_3 = 6.5, c_4 = 0.8$). To approximate this effective distance, the radius of influence in the potential field model was increased to $\rho_o = 8$ m (Fig. 18(a)). In the simulation (Fig. 18(b)), the agent starts in upper left corner with a heading of -28 degrees, almost facing the obstacle. The dynamical model generates the smooth trajectory labeled path 2. The potential field method traverses path 1(a) toward the obstacle, abruptly turns left, and continues on path 1(b). The forces acting on the depicted position of the agent are also represented, including the repulsive force of the obstacle, the attractive force of the target, and the resultant force vector. The obstacle repulsive force has a large longitudinal component but

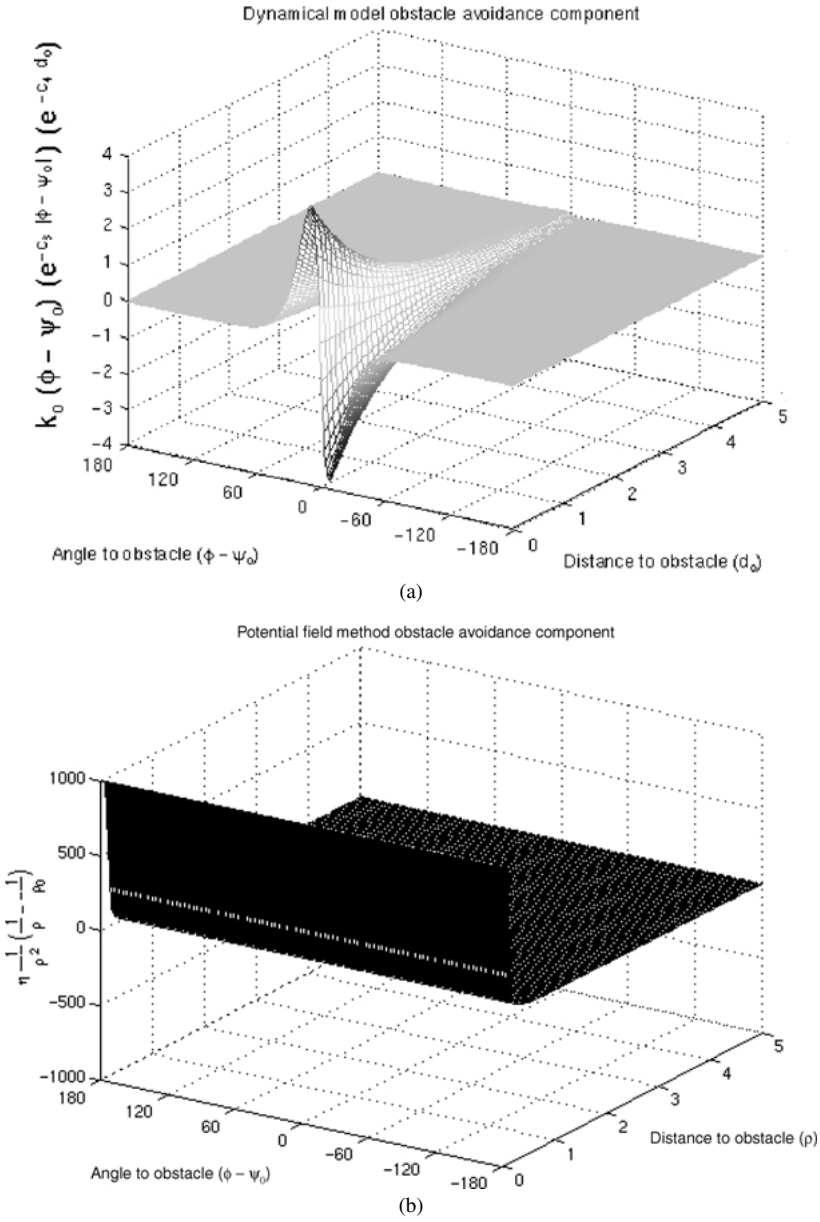

Figure 17. Obstacle avoidance components: (a) dynamical model and (b) potential field method. Parameter values for the dynamical model are $k_o = 66$, $c_3 = 6.5$, $c_4 = 0.8$, and for the potential field method are $\eta = 1$, $\rho_o = 4$.

a small lateral component (due to the nearly head-on approach to the obstacle), so the agent comes close to the obstacle before the lateral repulsive force steers it away.

In the second experiment, a smaller effective distance of repulsion (0.8 m) was selected for both methods (Fig. 19). The dynamical model generates path 2 and path 3 ($c_3 = 1$ and $c_3 = 0.5$, respectively, with $c_4 = 12$), with very close approaches to the obstacle and sharper turns. These paths demonstrate the importance of tuning both the distance and heading angle terms in the model's obstacle component together. With these parameter values, the effective distance is too short for the dynamical model, generating collision trajectories if the width of the agent is taken into account. The potential field method produces Path 1; the agent is depicted on this path with the corresponding force vectors.

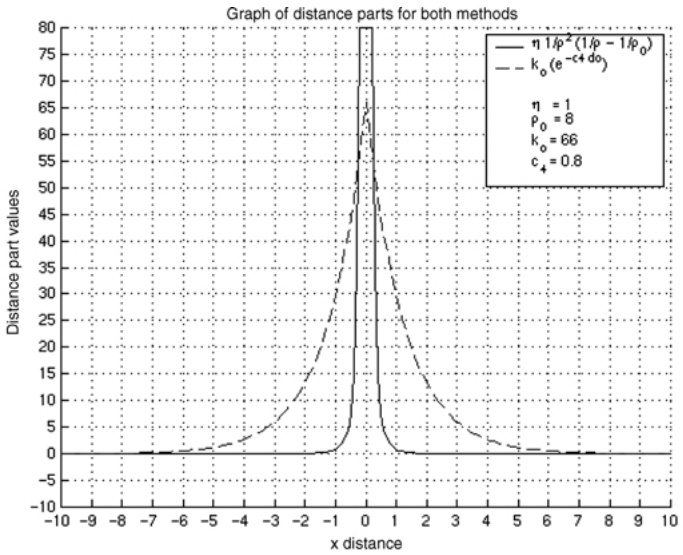

(a)

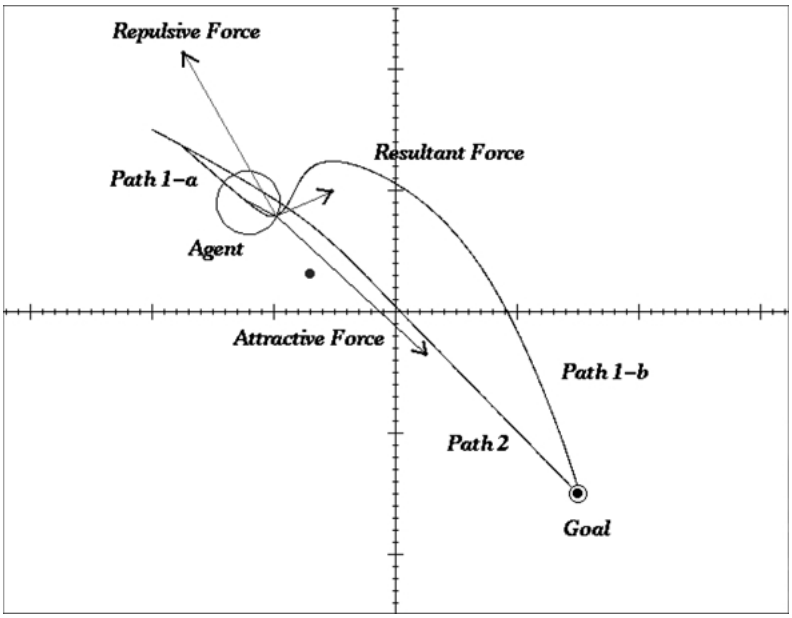

(b)

Figure 18. Experiment 1: (a) distance parts of both methods and (b) simulation results.

Fat Agents and Wide Obstacles. The potential field method inherently takes account of agent and obstacle width, because distance is measured between the boundary or envelope of the agent and that of the obstacle. In contrast, the current version of the dynamical model treats the agent and obstacles as points, and thus does not incorporate an explicit concept of width. Humans are very sensitive to the width of openings relative to their body size (Warren and Whang, 1987). The dynamical model implicitly expresses this relationship in the rate of exponential decay with obstacle distance (c_4 parameter). As illustrated in the previous section (Figs. 18 and 19), this determines how wide a berth the agent gives to an obstacle, and can thus be adjusted for body size.

However, the model is not yet designed to deal with wide obstacles. One possibility is simply to include the size of each obstacle as a parameter, but in a

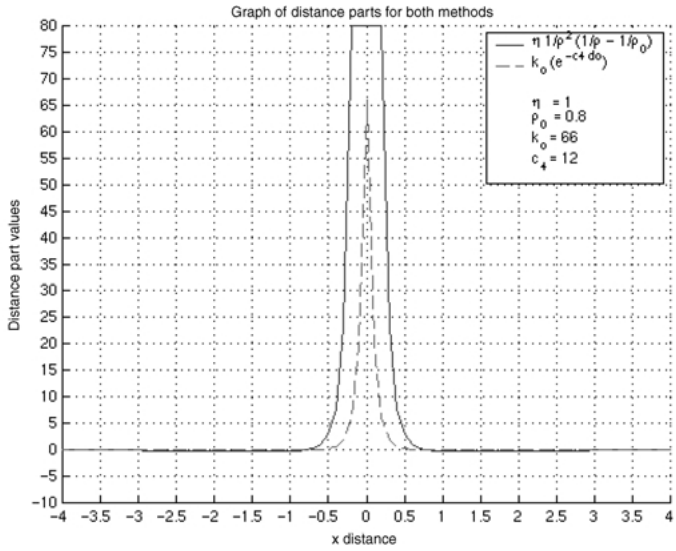

(a)

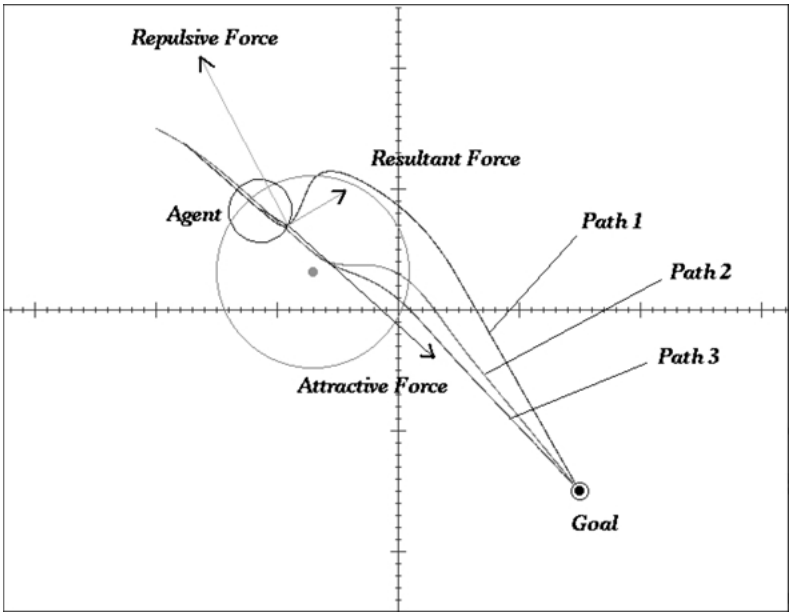

(b)

Figure 19. Experiment 2: (a) distance parts of both methods and (b) simulation results.

biologically-inspired model we seek to define the input in "proximal" terms such as visual angle rather than "distal" terms such as object size. Another approach would be to convolve the obstacle angle function over space, so that the entire visual angle of the obstacle is repulsive, rather than just a point at its center. It might also be possible to scale the obstacle angle function to the visual angle of the obstacle, which would cause nearby or large obstacles to be weighted more heavily.

Local Minima and Cancellation. The form of the obstacle function creates another important difference between the two approaches. In potential field methods, the magnitude of the repulsive force tends to infinity as the agent approaches the obstacle. This guarantees that the agent will never run over an obstacle. In the dynamical model, on the other hand, the obstacle influence is based on exponential decay and never produces infinite angular acceleration—a more realistic choice

for physical agents and humans. Combined with the difference in control variables (translational velocity vs. angular acceleration), this results in a significant advantage for the dynamical model, although it also creates a minor disadvantage.

Advantage. The potential field approach is a local obstacle avoidance method, and local minima are a serious problem. An agent using the potential field method alone without a high level path planner can easily get stuck in local minima, even in the simplest scenes. The dynamical model, in contrast, has few such problems, at least in simple scenes. Because it only controls angular acceleration and not the agent's speed (never stopping the agent), local minima are avoided in two ways: the agent either takes advantage of the *canceling effect* (described below) and passes between the obstacles (if the distance decay parameter c_4 is big), or it takes a path around the obstacle cluster (if c_4 is small). In the latter case it may overshoot the target, but it easily homes in from another direction. Thus, with appropriate parameter settings the dynamical model can avoid local minima in simple scenes.

Disadvantage. However, if the locations of the obstacles are symmetrical about the agent's path to the target, then their contributions to the angular acceleration will have similar magnitudes but opposite signs, and therefore cancel each other. This *canceling effect* creates a spurious attractor in the center of the obstacle array, which may lead the agent into a gap that is too small, or even to crash into an obstacle at the center of a perfectly symmetrical array. As noted above, one way to avoid the canceling effect is to increase obstacle repulsion with distance by reducing the exponential decay term c_4, thereby inducing an outside path around the entire array. In cases with only a few obstacles, adding a noise term to the model may allow it to escape unstable fixed points.

These advantages and disadvantages are illustrated in Fig. 20. In this example the agent starts in the lower left corner with an initial heading of 0°, and moves at a constant translation speed of 1 m/s. Path 1 shows a sample local minimum for the potential field method. The agent is stuck in a *bowl* (a region of small outward-pointing resultant vectors surrounded by large inward-pointing vectors) and is reduced to oscillating back and forth. Another type of local minimum is being *frozen* in a location where the attractive and repulsive forces cancel each other, producing a resultant force of zero magnitude. Path 2 is traversed with the dynamical model ($c_4 = 1.6$). Since there are obstacles on both sides of the agent, their combined contribution to the angular acceleration demonstrates the canceling effect along the path, and the agent passes between them. Path 3 is also traversed by the dynamical model using a more gradual exponential decay with distance ($c_4 = 0.4$). The repulsive regions of the obstacles are larger, and therefore they force the agent to take an outside path.

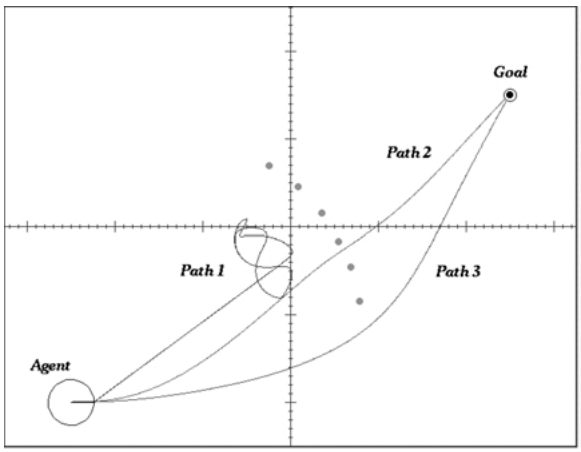

Figure 20. Example of a local minimum, canceling effect and outside path.

Agent Speed. A final difference between the two methods is that the dynamical model assumes a constant translational speed on the part of the agent. This is indeed the case in our human data: subjects tend to accelerate from a standstill and then maintain an approximately constant walking speed. However, the model produces different paths at different constant speeds, with all other parameters fixed. The reason for this behavior is that, when the agent enters a region that produces a non-zero angular acceleration, the accelerating effect lasts for a shorter time at higher speeds, inducing a smaller rotation. In contrast, since the potential field equations determine the direction of the agent's motion, it will always traverse the same path independent of speed. For any physical agent with mass and momentum, the responsiveness of trajectories to speed may actually be a desirable effect.

An example for the dynamical model is presented in Fig. 21. With a constant speed of 0.25 m/s, the model traverses path 1 to the left of the obstacle, but with a speed of 1.0 m/s it takes path 2 to the right. In these simulations, the agent's initial heading was 0° (horizontal),

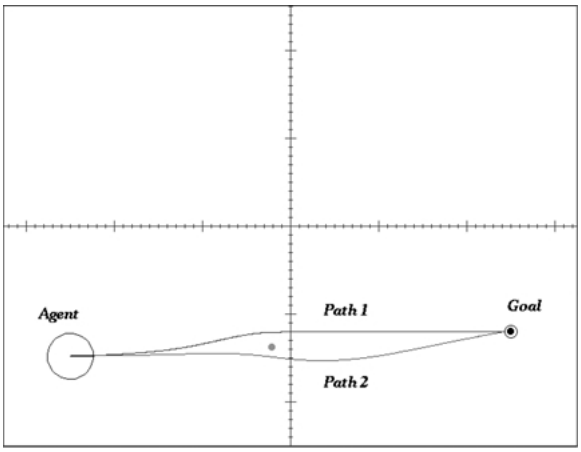

Figure 21. The effect of different translation speeds on the dynamical model generated paths.

as shown. On path 1, the agent initially rotated toward the target (counter-clockwise) and made a large enough turn before the obstacle's influence took effect that the obstacle pushed it to the left. On path 2, it began to rotate in the same direction (counter-clockwise), but because it was translating faster it did not have time to rotate past the obstacle before feeling its influence, so the obstacle pushed it to the right. It seems likely that walking speed will similarly affect human paths, but this remains to be empirically tested.

Although the human subjects in our experiments tended to maintain an approximately constant walking speed, there are likely to be certain situations in which humans decelerate if they get too close to an obstacle in order to avoid an immanent collision. Duchon et al. (1998) designed a robot that avoids collisions (in part) by decelerating when the robot's time-to-contact (Lee, 1976) with an obstacle reaches a margin value. A similar mechanism could easily be added to the dynamical model to guarantee collision avoidance.

Discussion

The aim of this project was to construct a dynamical model of steering, obstacle avoidance, and route selection that exhibits the stability and flexibility evident in human locomotion without relying upon an internal model of the environment. In this regard, we believe the model succeeds in capturing the behavioral dynamics. The human data were reproduced with fits near 1.0 and trajectories through the space of behavioral variables that closely resemble those produced by humans. In addition, it was shown how goal and obstacle components could be additively combined to yield smooth, stable trajectories in more complex scenes. The model thus scales linearly with the complexity of the scene. In practice, the model is even more computationally efficient, because steering can be based on a limited sample of a few meters ahead and close to the current heading direction, and does not require a memory representation of the complete 3D scene. The results demonstrate that on-line, information-based control is sufficient for steering, obstacle avoidance, and route selection in simple environments, without an explicit world model or advanced path planning. In effect, a route through the environment emerges from the agent's local responses to goals and obstacles.

Behavioral Dynamics and Laws of Control

In the introduction, we suggested that the behavioral dynamics arise from the implementation of laws of control in a physical agent that interacts with a physical environment. The present model provides a description of the behavioral dynamics, but it also allows us to make some inferences about laws of control. Because a physical agent is a 2nd-order dynamical system, the model of observed behavior is necessarily (at least) 2nd-order. However, following Schöner et al. (1995), it is advantageous for the control law to be 1st-order, so that the stability of its solutions can be assured. Thus, we can split the model into a 2nd-order physical system and a 1st-order control law. In the control law, angular velocity $\dot{\phi}$ is a function of the goal and obstacle components and quickly relaxes to an attractor for the desired heading, given sensory information about the current goal and obstacle angles. The physical system then determines the angular acceleration $\ddot{\phi}$ based on the difference between the desired and current heading, and a fixed damping term. Preliminary simulations have shown that this system produces comparable fits to the human data (with different parameter values), and thus gives rise to the behavioral dynamics described by the present model.

What visual information might specify the critical variable $\phi - \psi$, the angle of a goal or obstacle from the current heading direction? An obvious possibility is the angle between the egocentric direction of the object and the agent's locomotor axis, which could be determined from camera input or sonar sensors and knowledge of the orientation of the effectors (wheels).

Another possibility is the visual angle between the object and the heading specified by optic flow. There are a number of approaches to recovering heading from flow (Hildreth and Royden, 1998; Warren, 1998a), and this approach has the advantage that the control information is all in the visual domain. Recent evidence indicates that humans actually use both sources of information: egocentric direction dominates under low-flow conditions, but flow dominates as detectable motion is added to the scene (Rushton et al., 1998; Warren et al., 2001; Wood et al., 2000). The linear combination of redundant information about the angle $\phi - \psi$ provides for robust locomotor control under varying environmental conditions. Finally, note that even though we use the current distance d of an object as a variable in the model, absolute distance need not be recovered. An equivalent solution is to use the current *time-to-contact* with the object, approximated by the ratio of its visual angle to its rate of expansion (Lee, 1980; Tresilian, 1990), which can also be determined from the optic flow (Duchon et al., 1998).

We have emphasized how the information-based approach promotes the view that behavior is guided by occurrent information, rather than internal representations of the environment. Another type of internal representation often employed in mobile robots is an explicit model of the relationship between control variables and resulting body kinematics, that is, a model of the *plant dynamics*. Such a model is thought to be necessary for the agent to predict future states of the body. If behavior is a consequence of laws of control and physical constraints, however, then the agent does not need an explicit model of the plant dynamics. Rather, the agent learns a set of control law parameters that yield successful behavior for the given set of physical constraints. If these constraints change (e.g., if the agent's mass increases, or the medium changes from air to water), the agent may adapt by tuning the parameters until behavior is stabilized again. Thus, the agent's "model" of the plant dynamics is simply a set of parameter values that result in successful behavior within a given set of constraints.

Comparison with Other Approaches

The present model was inspired by the dynamical approach of Schöner et al. (1995) and thus has close affinities with their model. However, we depart from it in two related ways. First, the present model is intended as a description of the behavioral dynamics, spanning the agent's control laws and the physics of the agent and environment, whereas Schöner et al.'s (1995) model is intended as a control algorithm that treats the physics as an implementation detail. We believe that physical and biomechanical constraints codetermine behavior and may actually contribute to a solution, and thus should be incorporated in the model. Second, it follows that our model controls angular acceleration ($\ddot{\phi}$) rather than angular velocity ($\dot{\phi}$). This was motivated directly from measurements of human walking, and is a consequence of the fact that any physical agent has mass. Acceleration control contributes to the smooth, efficient paths exhibited by the model.

Detailed comparisons with the potential field approach revealed that the dynamical model has certain advantages. These include smoother, shorter trajectories, successful evasion of local minima in simple situations, and responsiveness to different speeds of travel. On the other hand, the model has yet to incorporate a concept of obstacle width and is subject to the canceling effect in symmetrical configurations. These seem to be tractable problems that may be dealt with by modifications to the model.

Conclusion

In sum, the present model provides a compact description of the behavioral dynamics of steering and obstacle avoidance that gives rise to the pattern of route selection exhibited by human subjects in simple scenes. Simulations of the model in more complex scenes, in which goal and obstacle components are linearly combined, revealed that the agent takes smooth, efficient paths to the goal. This suggests that route selection in autonomous agents need not require explicit path planning, but may emerge on-line as a consequence of the local steering dynamics.

Appendix

The full model is given by the following equation:

$$\ddot{\phi} = -b\dot{\phi} - k_g(\phi - \psi_g)(e^{-c_1 d_g} + c_2) + \sum_{i=1}^{\#obstacles} k_o(\phi - \psi_{o_i})e^{-c_3|\phi - \psi_{o_i}|}(e^{-c_4 d_{o_i}}) \quad \text{(A1)}$$

Note that ψ_g, d_g, ψ_o, and d_o change as the position of the agent changes (see Fig. 1). However, each of these variables can be expressed as a function of the (x, z)

position of the observer:

$$\psi_g = \cos^{-1}\left(\frac{(Z_g - z)}{d_g}\right) \tag{A2}$$

$$d_g = [(X_g - x)^2 + (Z_g - z)^2]^{1/2} \tag{A3}$$

$$\psi_o = \cos^{-1}\left(\frac{(Z_o - z)}{d_o}\right) \tag{A4}$$

$$d_o = [(X_o - x)^2 + (Z_o - z)^2]^{1/2} \tag{A5}$$

where (X_g, Z_g) and (X_o, Z_o) are the coordinates of the goal and obstacle, respectively. Written as a system of first-order differential equations, the model is given by:

$$\begin{aligned}
\dot{y}_1 &= \dot{\phi} \\
\dot{y}_2 &= \ddot{y}_1 = \ddot{\phi} = -b\dot{\phi} - k_g(\phi - \psi_g)(e^{-c_3 d_g} + c_2) \\
&\quad + \sum_{i=1}^{\#obstacles} k_o\big(\phi - \psi_{o_i}\big)e^{-c_3|\phi - \psi_{o_i}|}(e^{-c_4 d_{o_i}}) \\
\dot{y}_3 &= \dot{x} = V\sin\phi \\
\dot{y}_4 &= \dot{z} = V\cos\phi
\end{aligned} \tag{A6}$$

where V is the speed of the observer, which was held constant at 1 m/s in our simulations.

Acknowledgments

This research was supported by the National Eye Institute (EY10923), National Institute of Mental Health (K02 MH01353) and the National Science Foundation (NSF 9720327).

Note

1. Schöner et al. (1995) used the term "behavioral dynamics" to refer to a control algorithm independent of the physical implementation, which is closer to our notion of a "law of control." In contrast, we develop the concept of behavioral dynamics as a description of the observed behavior of the physical agent.

References

Aloimonos, Y. (Ed.) 1993. *Active Perception*. Erlbaum, Hillsdale, NJ.

Arkin, R.C. 1989. Motor schema-based mobile robot navigation. *International Journal of Robotics Research*, Aug:92–112.

Beer, R. 1997. The dynamics of adaptive behavior: A research program. *Robotics and Autonomous Systems*, 20:257–289.

Beer, R.D. 1995. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215.

Brooks, R.A. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2:12–23.

Brooks, R.A. 1991. Intelligence without representation. *Artificial Intelligence*, 47:139–160.

Duchon, A.P., Warren, W.H., and Kaelbling, L.P. 1998. Ecological robotics. *Adaptive Behavior*, 6:473–507.

Fajen, B. and Warren, W.H. 2003. The behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2):343–362.

Franceschini, N., Pichon, J.M., and Blanes, C. 1992. From insect vision to robot vision. *Philosophical Transactions of the Royal Society of London, B*, 337:283–294.

Gibson, J.J. 1958/1998. Visually controlled locomotion and visual orientation in animals. *British Journal of Psychology*, 49:182–194. Reprinted in *Ecological Psychology*, 10:161–176.

Gibson, J.J. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston.

Hildreth, E.C. and Royden, C.S. 1998. Computing observer motion from optical flow. In *High-Level Motion Processing*. T. Watanabe (Ed.), MIT Press, Cambridge: MA, pp. 269–293.

Kelso, J.A.S. 1995. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. MIT Press, Cambridge: MA.

Khatib, O. 1986. Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5:90–98.

Krogh, B.H. 1984. A generalized potential field approach to obstacle avoidance control. *International Robotics Research Conference*, Bethlehem, PA.

Kugler, P.N. and Turvey, M.T. 1987. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*. Erlbaum, Hillsdale, NJ.

Large, E.W., Christensen, H.I., and Bajcsy, R. 1999. Scaling the dynamic approach to path planning and control: Competition among behavioral constraints. *International Journal of Robotics Research*, 18:37–58.

Lee, D.N. 1980. Visuo-motor coordination in space-time. In *Tutorials in Motor Behavior*. G.E. Stelmach and J. Requin (Eds.), North-Holland, Amsterdam, pp. 281–295.

Moravec, H.P. 1981. Obstacle avoidance and navigation in the real world by a seeing robot rover. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, William Kaufmann, Los Altos, CA, pp. 785–790.

Newman, W.S. and Hogan, N. 1987. High speed robot control and obstacle avoidance using dynamic potential functions. In *Proceedings of the 1987 IEEE International Conference on Robotics and Automation*, Raleigh, NC, pp. 14–24.

Pfeiffer, F., Eltze, J., and Weidemann, H.J. 1994. The TUM walking machine. In *Intelligent Automation and Soft Computing*. M. Jamshidi, C. Neguyen, R. Lumia, and J. Yuh (Eds.), TSI Press: Albuquerque, NM.

Ritzmann, R.E., Quinn, R.D., Watson, J.T., and Zill, S.N., 2000. Insect walking and biorobotics: A relationship with mutual benefits. *Bioscience*, 50:23–33.

Rushton, S.K., Harris, J.M., Lloyd, M., and Wann, J.P., 1998. Guidance of locomotion on foot uses perceived target location rather than optic flow. *Current Biology*, 8:1191–1194.

Schöner, G. and Dose, M., 1992. A dynamical systems approach to task-level system integration used to plan and control autonomous vehicle motion. *Robotics and Autonomous Systems*, 10:253–267.

Schöner, G., Dose, M., and Engels, C., 1995. Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16:213–245.

Srinivasan, M.V. and Venkatesh, S. (Ed.), 1997. *From Living Eyes to Seeing Machines*. Oxford University Press: Oxford.

Strogatz, S.H. 1994. *Nonlinear Dynamics and Chaos*. Addison-Wesley: Reading, MA.

Temizer, S. 2001. Optical flow based local navigation. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Temizer, S. and Kaelbling, L.P. 2001. Holonomic planar motion from non-holonomic driving mechanisms: The front-point method. In *Proceedings of Photonics Boston, Intelligent Systems and Advanced Manufacturing (Mobile Robots XVI)*, Newton, MA.

Thorpe, C.F. 1985. Path relaxation: Path planning for a mobile robot, Carnegie-Mellon University, The Robotics Institute, Mobile Robots Laboratory, Autonomous Mobile Robots, Pittsburgh, PA.

Tresilian, J.R. 1990. Perceptual information for the timing of interceptive action. *Perception*, 19:223–239.

Warren, W.H. 1988. Action modes and laws of control for the visual guidance of action. In *Movement Behavior: The Motor-Action Controversy*. O. Meijer and K. Roth (Eds.), North Holland, Amsterdam, pp. 339–380.

Warren, W.H. 1998a. The state of flow. In *High-Level Motion Processing*. T. Watanabe (Ed.), MIT Press, Cambridge, pp. 315–358.

Warren, W.H. 1998b. Visually controlled locomotion: 40 years later. *Ecological Psychology*, 10:177–219.

Warren, W.H., Kay, B.A., Zosh, W.D., Duchon, A.P., and Sahuc, S. 2001. Optic flow is used to control human walking. *Nature Neuroscience*, 4:213–216.

Warren, W.H. and Whang, S., 1987. Visual guidance of walking through apertures: Body scaled information for affordances. *Journal of Experimental Psychology: Human Perception and Performance*, 13:371–383.

Wood, R.M., Harvey, M.A., Young, C.E., Beedie, A., and Wilson, T. 2000. Weighting to go with the flow? *Current Biology*, 10:R545–R546.