Perception, 2005, volume 34, pages 717 – 740

DOI:10.1068/p5405

Perceiving possibilities for action: On the necessity of calibration and perceptual learning for the visual guidance of action

Brett R Fajen
Department of Cognitive Science, Rensselaer Polytechnic Institute, Carnegie Building 308, 110 Eighth Street, Troy, NY 12180-3590, USA; e-mail: fajenb@rpi.edu
Received 15 November 2003, in revised form 7 June 2004; published online 15 June 2005

Abstract. Tasks such as steering, braking, and intercepting moving objects constitute a class of behaviors, known as visually guided actions, which are typically carried out under continuous control on the basis of visual information. Several decades of research on visually guided action have resulted in an inventory of control laws that describe for each task how information about the sufficiency of one's current state is used to make ongoing adjustments. Although a considerable amount of important research has been generated within this framework, several aspects of these tasks that are essential for successful performance cannot be captured. The purpose of this paper is to provide an overview of the existing framework, discuss its limitations, and introduce a new framework that emphasizes the necessity of calibration and perceptual learning. Within the proposed framework, successful human performance on these tasks is a matter of learning to detect and calibrate optical information about the boundaries that separate possible from impossible actions. This resolves a long-lasting incompatibility between theories of visually guided action and the concept of an affordance. The implications of adopting this framework for the design of experiments and models of visually guided action are discussed.

1 Introduction

The prospective control of action is fundamental to the successful performance of both routine, everyday acts, such as navigating through a busy shopping mall, and highly skilled tasks, such as positioning oneself to catch a fly ball in baseball. In both cases, success depends on the ability to perceive the behavioral possibilities of the environment, so that one's movements can be appropriately coordinated (Turvey 1992). When an animal unexpectedly darts out in front of an automobile, the driver must know whether or not it is possible to slow down quickly enough to avoid a collision. If the animal is too close or the automobile is moving too quickly to stop in time, the driver may opt to swerve out of the way. Likewise, a competitive cyclist, when approaching a sharp bend in the road, must know whether or not it is possible to turn quickly enough to stay on the road without losing traction. If the bend is too sharp or the surface is wet, then the cyclist may slow down before entering the curve.

These tasks constitute a class of behaviors, which are referred to as visually guided actions because they are typically carried out under continuous control on the basis of visual information. Several decades of research on visually guided action have resulted in an inventory of control laws that describe for each task how observers are informed about the sufficiency of their current state (assuming conditions persist), and how that information is used to make ongoing adjustments. Although this framework for studying visually guided action has proven to be successful in many respects, several aspects of these tasks that are essential to successful performance cannot be captured. The purpose of this paper is to provide a novel way of looking at the capabilities that underlie successful human performance in these tasks, and, it is hoped, encourage a new round of investigation that will lead to a deeper understanding of the visual guidance of action.

This paper begins with a look at the historical context within which research on visually guided action evolved. Existing models of steering, braking, interception, and fly-ball catching are then reviewed to illustrate the common assumptions that underlie research on all of these tasks. The inadequacies of these models are described, and a new framework that emphasizes the significance of calibration and perceptual learning is introduced. It is then argued that sensitivity to the boundaries that separate possible from impossible actions is the fundamental basis of successful human performance in these tasks, which resolves a long-lasting incompatibility between theories of visually guided action and the concept of an affordance (Gibson 1986). Lastly, I consider the implications of adopting this framework for designing experiments and models of visually guided actions.

2 Historical context

In 1958, James J Gibson published a landmark paper titled "Visually controlled locomotion and visual orientation in animals". This paper is regarded as seminal for several reasons.[(1)] First, Gibson provided a convincing justification for the scientific study of visually guided locomotion by those interested in visual space perception. As Warren (1998) put it, "the implication was that perceiving the layout of surfaces and what it affords for locomotor behavior is the evolutionarily primary function of spatial perception" (page 182). Second, Gibson proposed a set of formulae for routine locomotor tasks, such as steering toward a goal, turning to avoid an obstacle, slowing down to avoid an obstacle, and intercepting a moving target. These formulae continue to provide the inspiration for a great deal of empirical research on visually guided action. Third, in his paper, Gibson showed how visually guided locomotion could, in principle, emerge naturally from the detection of information without the need for a detailed internal model of the environment. These ideas anticipated much of the current research on steering, obstacle avoidance, and navigation in humans (Fajen and Warren 2003; Fajen et al 2003), insects (Srinivasan 1998), and robots (Duchon et al 1998).

Much of the substantial body of current research on visually guided action is aimed at testing, refining, and formalizing Gibson's intuitions (see Warren 1998 for a review). In addition, control strategies for several new tasks, such as catching a fly ball (Chapman 1968; McBeath et al 1995; Michaels and Oudejans 1992), maintaining altitude (Flach et al 1997), and following a lane (Beall and Loomis 1996; Duchon and Warren 2002), have been proposed and tested. When expressed formally, as they often are in the current literature, these strategies describe how information variables map onto action variables. Out of necessity, different tasks involve different sets of informational and action variables, and so the details of each model differ. Yet there are striking similarities among most current models of visually guided action that reflect often implicit assumptions about how information is used to control action.

2.1 *The error-nulling assumption*

The first assumption is that performing a visually guided action involves some form of error nulling. For each task, one could define an 'ideal state' in which the observer should strive to be at each moment in time. When steering to a goal, for example, the ideal state is the path curvature that would bring the observer to the goal without making any further steering-wheel adjustments along the way. When braking, the ideal state is the rate of deceleration that would bring the observer to a stop at the desired location without making any further brake adjustments. The difference between the current state and the ideal state constitutes an error that must be corrected (or 'nulled') for the action to be successful. If the current and ideal states could be estimated, then error-nulling adjustments could be based on their difference. However, it is generally agreed that such states are not always perceptually available to the observer.

[(1)] Its significance was recently acknowledged in a special issue of the journal *Ecological Psychology* (1998 **10** 3 – 4) marking the 40th anniversary of the publication of the paper.

The solution that has been adopted in the literature can be traced back to Gibson's attempts to understand how movements can be controlled in the presence of external forces, such as the flow of wind or water. When such forces are present, the observer's motion relative to the environment and the resulting pattern of optic flow are due to a combination of externally and internally generated forces (Warren 1988). To know how much internal force to apply, it would seem that the observer must know and compensate for the presence of external forces. If the observer could estimate and subtract the component of optic flow due to internal forces on the basis of an efferent copy of the motor command (von Holst 1954), then the remaining component could be attributed to external forces and used to estimate the required compensation. The problem is that the relation between motor commands and optical consequences is complex and usually unpredictable owing to the presence of context-conditioned variability (Bernstein 1967; Turvey et al 1982). Gibson's (1986) solution was to rely on optical information about one's movements relative to the environment, which he called 'visual kinesthesis'. In order for movements to be controlled in the presence of external disturbances such as crosswinds and gravity, the observer must act so as to achieve a certain type of optic flow. To illustrate Gibson's solution, Warren (1988) used the simple but instructive example of an organism intending to maintain its position in the presence of external disturbances. Because a global transformation of the optic array corresponds to movement of the observer through a stable environment, the ideal state of affairs (ie zero velocity) can be achieved by applying forces to cancel any change in the optic array. For example, if a sudden gust of wind pushes the organism from its current position, the resulting global transformation of the optic array specifies the change in force that must be produced by the organism to restore the ideal state of affairs. Thus, regardless of the origin of the forces (internal or external) that displace the organism from its position, acting so as to cancel the change in the optic array will bring about the ideal state.

For these principles to be generalized to more complex visually guided actions, it must be shown that the ideal state can be attainable by acting so as to achieve a certain pattern of optic flow. In the case of fly-ball catching, for example, it must be shown that the ideal running speed can be attained by acting so as to produce a certain pattern of flow. As long as the fielder is producing the right flow pattern, the current and ideal running speeds are the same and no further adjustments are necessary. When the flow pattern is not what it should be, the current and ideal states are unequal and adjustments must be made if the fly ball is to be caught. Thus, information contained in optic-flow fields specifies the sufficiency of the observer's current state. By detecting this information, the observer is able to stay informed about his or her potential future, given current conditions, so that corrective adjustments can be made. This provides a simple, smart solution to the problem of how one nulls the error—not by estimating the current and ideal states, but by acting so as to produce a certain flow pattern. The challenge is to discover the flow pattern for each task.[(2)]

To summarize, the current perspective has compelled researchers to think of visually guided action as an 'error-nulling' process. The difference between the current and ideal states constitutes an error that must be corrected if the action is to be successful. Information about the sufficiency of one's current state functions as the input to the system, and the law of control describes how such information is used to make corrections to cancel the error.

[(2)] In the simple case of maintaining position, the ideal state is specified by a simple property of the optic-flow field (ie no global change in the optic array). For more complex actions with ideal states constituted by higher-order motions of the observer, the corresponding properties of the optic-flow field must be higher-order as well.

2.2 *The single-optical-invariant assumption*

A second assumption found in the literature on visually guided action is that successful performance across a variety of conditions can be attributed to the use of a single optical variable for each task that specifies the relevant property of the actor – environment system—that is, an optical invariant. The basis of this assumption can be traced back to several principles promoted by Gibson (1966, 1986) and has been embraced by those who subscribe to his ecological approach (Michaels and Carello 1981; Turvey et al 1981). These principles help guide decisions about what one is willing to count as information. In the ecological view, perception is lawfully related to the environmental properties. Because patterns in ambient-energy arrays link the perceiver and the environment, the ecological view requires that perception be related one-to-one to ambient-energy patterns, and that such patterns (appropriately described) be related one-to-one to the perceived properties. The term 'information' is therefore reserved for ambient-energy patterns that contain a specification of those properties. This view has the advantage that it keeps the animal in contact with the relevant properties of the environment without the need for mental entities that mediate this coupling.

The search for specification focuses on invariants, which are patterns of ambient-energy arrays that are left unchanged by certain transformations. For example, the optical variable τ (Lee 1976), equal to the inverse of the relative rate of optical expansion of an approaching object, is often cited as an example of an optical invariant because it uniquely specifies one's time-to-contact across changes in object size. In contrast, lower-order optical variables, such as optical angle and expansion rate, are affected by factors such as object size. Thus, whereas τ is said to be a specifying variable, optical angle and expansion rate are nonspecifying variables.[3] Because specifying variables are invariant across changes in conditions that are typically encountered, they support successful performance across a variety of conditions. Nonspecifying variables may work 'well enough' in constrained environments, but do not lawfully relate to any particular property of the environment (see Jacobs and Michaels 2002, for more on this issue). Thus, empirical research motivated by these principles is aimed at identifying optical variables that specify relevant environmental properties, and designing experiments to confirm that people actually rely on such variables to perceive the environment and control their actions.

The single-optical-invariant assumption and the error-nulling assumption go hand-in-hand. Recall that the ideal state for a visually guided action is attained by acting so as to bring about a certain pattern of optic flow, and that such flow patterns are characterized by some higher-order property that is left unchanged (ie invariant) across conditions whenever the observer is in the ideal state. Thus, the search for laws of control is (in part) the search for an optical invariant.

To summarize, theories and models of visually guided action are shaped by two widely accepted assumptions:

(i) The error-nulling assumption: observers make adjustments to null the difference between the current and ideal states by relying on information about the sufficiency of their current state.

(ii) The single-optical-invariant assumption: observers rely on a single optical variable for each task that is invariant whenever the observer is in the ideal state.

In the next section, existing models of visually guided action are reviewed to illustrate how they conform to this approach.

[3] Nonspecifying variables have also been referred to as 'incomplete invariants' (Runeson and Vedeler 1993) and 'perceptual heuristics' (Gilden and Proffitt 1989, 1994). The term 'nonspecifying variable' is used in this paper because it is more theoretically neutral than the other two.

3 Models of visually guided actions

3.1 *Steering toward a stationary goal*

To walk to a stationary goal, observers could simply align the locomotor axis with the goal using global optic flow, retinal flow, or egocentric direction (Rushton et al 1998; Warren et al 2001), and walk straight ahead. Although this simple strategy may work for a walking observer capable of nearly instantaneous changes in direction, it does not apply to the more general case in which turning rate is constrained by the controller dynamics (eg a steering wheel or handlebars) or by momentum, forcing the observer to follow a curvilinear trajectory to the goal. A number of possible strategies for steering a curved path to a stationary goal have been proposed. Lee and Lishman (1977) noted that when an observer follows a circular path over a textured ground plane, the streamline in the optic-flow field that passes out of view directly beneath the observer (ie the locomotor flow line) specifies the path that he or she will follow if no steering adjustments are made. In fact, observers can perceive their future circular path on the basis of the global optic-flow field even when the locomotor flow line is absent (Warren et al 1991). This means that observers could correct the steering-wheel-angle error by aligning the perceived future path with the goal. When the perceived future path is to the outside (or inside) of the goal, the observer should increase (or decrease) steering-wheel angle.

Another steering model was proposed by Kim and Turvey (1999; see also Wann and Swapp 2000), who created computer simulations of the retinal-flow pattern produced by an observer fixating a target while steering. They found that, when an observer fixates a point along the future circular path, the flow lines in the raw retinal-flow field are linear. Oversteering and understeering produce flow lines that curve toward or away from the direction of the steering error, respectively.

Wann and Land (2000) suggested another steering strategy based on the visual direction of the goal with respect to the instantaneous direction of heading. If the curvature of the observer's path is such that he or she will reach the target by maintaining the current steering-wheel angle, then the rate of change of the heading angle will be constant. So, if the rate of change of the heading angle increases or decreases, then the observer is oversteering or understeering.

Although these models differ in terms of the specific optical invariant used to control steering, the common assumption is that the optical invariant must inform observers about the sufficiency of the current path curvature, and is used to make adjustments to correct or null the steering error between the current and ideal path curvature.

3.2 *Intercepting a moving target on foot*

A well-known rule of thumb used by pilots and sailors to avoid a collision with a moving obstacle is to adjust one's speed and direction so that the bearing angle of the obstacle does not remain constant. It follows that a plausible strategy for achieving, rather than avoiding, collision with a moving start is to move so as to maintain a constant bearing angle. If the bearing angle increases, the observer could slow down or turn in the direction of the goal. If the bearing angle decreases, the observer could speed up or turn ahead of the goal. Thus, the change in bearing angle specifies the sufficiency of one's current speed and direction, and the 'constant-bearing-angle' strategy predicts that observers will change speed or direction to null the change in bearing angle. Empirical support for the constant-bearing-angle strategy was provided by Fajen and Warren (2004), and Lenoir et al (1999).

3.3 *Catching a fly ball*

Chapman (1968) proposed a model of fly-ball catching according to which the fielder runs forward or backward so as to cancel the optical acceleration of the tangent of the angle α subtended by the ball and the horizon. When $\mathrm{d}^2 \tan\alpha/\mathrm{d}t^2$ is negative, then the ball will drop in front of the fielder if current running speed is maintained. To catch the ball, the fielder must increase forward running speed until $d^2 \tan\alpha/dt^2$ is equal to zero. Likewise, when $d^2 \tan\alpha/dt^2$ is positive, the ball will fly over the fielder unless he or she accelerates in the backwards direction. Thus, $d^2 \tan\alpha/dt^2$ provides information about the sufficiency of one's running speed, and predicts that fielders will move so as to cancel the optical acceleration of the ball. This has come to be known as the optical acceleration cancellation (OAC) model.

OAC can be used to account for how fielders run forward and backward to catch a fly ball hit directly at them, but not how they move laterally to catch balls hit to the side.[4] McBeath et al (1995) derived an alternative model of fly-ball catching, which they referred to as the linear optical trajectory (LOT) model, that provides a unified account of fielders' fore–aft and lateral movement. The LOT model predicts that fielders change both speed and direction so as to cancel the optical curvature (rather than the optical acceleration) of the ball. So the LOT strategy allows for more flexibility in the running path, curved trajectories, and changes in running speed.

The strategy that fielders actually use to catch fly balls is the focus of active, ongoing debate in the literature (McLeod et al 2001; Shaffer and McBeath 2002; Zaal and Michaels 2003). But despite the differences between the OAC and LOT models, a common assumption of both models is that fielders rely on information about the sufficiency of their current running speed and direction, and make adjustments to cancel the difference between the current and ideal running speed and direction.

3.4 *Intercepting a moving target by hand*

Models of fly-ball catching describe how fielders position themselves in the outfield to catch a fly ball, but not how they actually move the glove or hand to the right position to make the catch. Peper, Bootsma, and colleagues (Bootsma et al 1997; Peper et al 1994) focused on a simplified version of this problem that involves moving the hand along the frontal plane of the body to catch a ball as it passes by the observer. According to their 'required-velocity' model, the lateral acceleration of the hand is equal to the difference between the current hand velocity and the hand velocity required to intercept the ball. The required (or ideal) velocity is equal to the distance between the lateral position of the hand and the ball, divided by the time-to-contact of the ball with the frontal plane. Thus, the model describes how information about the sufficiency of one's current hand velocity is used to cancel the difference between the current and ideal hand velocity, which brings the hand into the right place at the right time to make a catch. Empirical support for this model was provided by Montagne et al (1999).

3.5 *Braking to avoid an obstacle in the path of motion*

As an observer approaches an object, the optical angle θ subtended by the contours of the object undergoes expansion, and the optical variable $\tau = \theta/\dot{\theta}$ decreases (at least throughout most of the approach). The first temporal derivative of τ ($\dot{\tau}$ or 'tau-dot') specifies the sufficiency of one's current deceleration (Lee 1976). When $\dot{\tau} < -0.5$, deceleration is too low, such that the observer will crash if brake pressure is maintained. When $\dot{\tau} > -0.5$, deceleration is sufficient, and the observer will stop short of the object if brake pressure is maintained. Thus, $\dot{\tau}$ functions as a control variable for visually guided braking. To stop as closely as possible to an object, the observer should increase brake pressure when $\dot{\tau} < -0.5$, and decrease brake pressure when $\dot{\tau} > -0.5$. This will result in adjusting the brake in the direction that nulls the difference between the current and ideal deceleration; that is, nulling the $\dot{\tau}$ error nulls the deceleration error.

[4] According to Chapman (1968), the fielder's lateral movements are controlled by maintaining a constant bearing angle in the same way that one intercepts a target moving along the ground. Thus, the fielder's fore–aft and lateral movements are controlled by two independent strategies.

Using an actively controlled braking task, Yilmaz and Warren (1995) provided evidence that $\dot{\tau}$ is actually used to control braking (see Bardy and Warren 1997 for a review).

4 A critique of existing models of visually guided actions

The preceding review of existing models of visually guided actions reveals that attempts to understand these tasks are aimed at describing how information about one's potential future (assuming current conditions persist) is used to correct deviations from the state of the action system required to achieve the task (ie the ideal state). In many respects, this has been an extremely successful enterprise. The more parsimonious 'information-based' approach has proven to be sufficient, paving the way for the rejection of the assumption that an internal 3-D model of the environment is necessary (Warren 1998). Furthermore, each of these models makes specific, empirically testable predictions about an ideal state in which the observer should attempt to be at each moment. When the data from individual trials are analyzed, observers often appear to be making adjustments to move in the direction of the ideal state (eg Yilmaz and Warren 1995).

4.1 *The problem of action boundaries*

Despite the successful, widespread application of this framework, there are important aspects of visually guided action that existing models fail to capture. Consider the situation in which the driver of an automobile is approaching a toll booth several hundred meters away at a typical highway speed. Even if the ideal deceleration is greater than the current deceleration, drivers do not typically initiate braking immediately upon recognition of the toll booth because doing so would lead to an unnecessarily slow approach. Rather, it is common to wait until the automobile is closer to the toll booth so that the driver can decelerate in a more time-efficient (but still safe and comfortable) manner. How long can the driver wait before initiating deceleration? The answer depends on several factors, including the strength of the brake, the traction of the vehicle, and the inertial forces that the driver is willing to tolerate. Even after braking is initiated, observers almost never exactly null the deceleration error. Overshoots and undershoots are common, as are periods in which brake pressure is held constant. The fact is that nulling the difference between the ideal and current deceleration throughout the entire approach is not a necessary condition for successful performance. More importantly, although it is a sufficient condition for success, it is not a condition that is always possible to satisfy. If there were no limits to the rate of deceleration that one could achieve, then all deceleration errors (no matter how large) could be nulled. However, such limits do exist, and they impose a critical constraint on successful braking. If the driver in the example above waits too long to start braking, or increases brake pressure too slowly, then at some point the ideal deceleration will exceed the maximum possible deceleration and the driver will be unable to null the deceleration error even by slamming on the brakes.

Thus, maximum deceleration defines an action boundary for braking that separates situations into two categories—those in which it is still possible to stop and those in which it is no longer possible to stop. Given the necessity of respecting this action boundary, observers must be extremely sensitive to the limits of their action capabilities. In fact, one might argue that such sensitivity is the key to successful braking. Unfortunately, existing models fail to capture this aspect of braking because the optical variable upon which observers presumably rely to control braking ($\dot{\tau}$) specifies the sufficiency of current deceleration, which is completely independent of maximum deceleration. If observers rely on $\dot{\tau}$ to control braking, then they would be unable to know whether the appropriate brake adjustment necessitates exceeding maximum deceleration.

The same principle applies to each visually guided action described in the previous section. When steering toward a distant goal, for example, the observer does not need to initiate turning immediately to match the ideal steering-wheel angle. Once turning is initiated, there is a tendency to overshoot the ideal steering-wheel angle and approach along a linear trajectory rather than a circular trajectory as one would expect from an idealized error-nulling model of steering. The advantage of approaching along a linear path is that small heading errors can easily be overcome by turning the steering wheel. If the observer approaches along a circular path close to the maximum turning rate, then it may be impossible to increase turning rate enough to overcome a heading error. This is analogous to the risk involved in keeping the ideal deceleration close to the maximum deceleration when braking. Generally speaking, for performance to be successful, the ideal state must never exceed the maximum capabilities of the observer.

4.2 *The problem of nonspecifying variables*

A second problem with the existing framework for visually guided action research arises from the assumption that successful performance across a variety of conditions can be attributed to the use of a single optical variable for each task—that is, the single-optical-invariant assumption. Although this assumption provides a parsimonious account of visually guided action, the possibility that people rely on nonspecifying variables to control steering, braking, interception, and other closed-loop actions has not been seriously considered. This is somewhat surprising, given the attention that this issue has received in studies on tasks such as catching, hitting, and making perceptual judgments about kinetic properties of the environment. Several interesting conclusions have emerged from these studies:

(i) People do rely on nonspecifying variables at least some of the time. Michaels et al (2001) reported that volleyball players use the optical expansion rate of the ball, rather than τ, to time the initiation of elbow flexion when punching a falling ball. Smith et al (2001) found a similar pattern of results for novices performing a task in which participants had to time the release of a pendulum to strike an approaching ball. In addition, in studies on one-handed catching (van der Kamp et al 1997), collision avoidance (DeLucia and Warren 1994), and time-to-contact estimation (Caird and Hancock 1994; DeLucia 1991) 'size effects' consistent with the use of optical angle or expansion rate rather than τ have been reported. In an entirely different domain, researchers have shown that perceptual judgments of kinetic properties of the environment, such as the peak force exerted by bimanual pullers (Michaels and de Vries 1998), and the relative mass of colliding balls (Jacobs et al 2001), are sometimes based on lower-order, nonspecifying variables.

(ii) Different individuals rely on different optical variables. For example, Jacobs et al (2001) and Michaels and de Vries (1998) found that some observers rely on specifying variables while other observers use nonspecifying variables to make judgments of relative mass and peak force.

(iii) Observers may converge on higher-order, specifying variables with practice and experience. Smith et al (2001) found that observers began to rely on a more 'τ-like' combination of angle and expansion rate after several sessions of practice on the pendulum-timing task. Some observers in their study eventually used τ. Similarly, Jacobs et al (2001) reported that observers asked to judge the relative mass of colliding balls switched to more effective variables with practice and feedback.

(iv) Attunement to one optical variable rather than another can depend on factors such as the range of conditions experienced by the observer, the presence of local constraints, viewing conditions (eg monocular versus binocular), and the criteria for success. Such factors can considerably change the effectiveness of a given optical variable. For example, in Smith et al (2001), one group of participants practiced the pendulum-timing task using a slow range of approach speeds and another group practiced using a fast range of approach speeds. Expansion rate was an effective variable in the fast condition, but not in the slow condition. Differences in performance between the two groups at overlapping speeds suggested that participants in the fast group relied on expansion rate, whereas those in the slow group learned to use a linear combination of angle and expansion rate. Smith et al concluded that the ability to perform a task across a wide range of conditions reflects a form of perceptual learning, which they referred to as 'flexible attunement' because observers are able to flexibly attune to different optical variables as conditions change. Similarly, Jacobs et al (2001) demonstrated that convergence onto higher-order variables can be facilitated by exposing observers to practice conditions in which the degree of correlation of commonly used nonspecifying variables with the to-be-perceived property is weak.

These findings are forcing many researchers to abandon the single-optical-invariant assumption, and instead search for the principles that explain when and why observers select certain optical variables and switch to others (Tresilian 1999). However, none of these studies concern the kinds of visually guided actions that are the focus of this paper. Within the literature on visually guided action, the focus is almost exclusively on optical invariants. Consequently, the study of skill acquisition and adaptation to changing conditions as attunement to more effective optical variables with experience has not been pursued in the literature on visually guided action.

Why do researchers who study visually guided action focus exclusively on optical invariants despite the overwhelming evidence indicating that observers sometimes rely on nonspecifying variables? If one begins with the error-nulling assumption, then it is almost impossible to imagine how observers would use anything other than an optical invariant. Consider how a nonspecifying variable might be used to control braking from the perspective that people attempt to null the deceleration error. Observers would still make brake adjustments around a critical value of some optical variable (eg optical expansion rate), but the optical variable would not specify the sufficiency of one's current deceleration as $\dot{\tau}$ does. Consequently, observers may adjust the brake in the wrong direction—that is, they may decrease brake pressure when they should increase it (or vice-versa) to null the deceleration error. Performance would be systematically biased, but that is what one would expect when nonspecifying variables are used. The problem that researchers encounter is that it is not clear how to make predictions about the specific critical value of a given nonspecifying variable. Calculating the predicted critical value of $\dot{\tau}$ (ie -0.5) was straightforward. But suppose a researcher wanted to test whether observers regulate braking around a critical value of expansion rate. Around what critical value would we expect observers to regulate braking? 0.5 rad s^{-1}? 1.0 rad s^{-1}? Different critical values lead to different levels of performance depending on the range of conditions, but there is no established method for measuring the effectiveness of a specific critical value of a nonspecifying variable. Furthermore, it is unclear how one would analyze the time-series data from experiments on braking to test whether observers regulate deceleration around the predicted critical value of a given nonspecifying variable (as opposed to a different critical value, or as opposed to the predicted critical value of another nonspecifying variable). In short, if one begins with the assumption that brake adjustments are made to null the deceleration error, then designing experiments to test for the possible use of nonspecifying variables is unfeasible. This explains why researchers in visually guided action focus exclusively on optical invariants despite the overwhelming evidence coming from other areas of research indicating that observers do not rely exclusively on specifying variables. An unfortunate consequence is that the role of perceptual learning as attunement to more effective optical variables with practice has not been investigated in the literature on visually guided action.

In sum, the theoretical framework within which existing research on visually guided action is conducted fails to account for at least two important aspects of successful performance. First, the maximum limits of one's action capabilities impose critical constraints on successful performance. Observers must be extremely sensitive to the limitations of their action capabilities, yet existing models cannot account for such sensitivity. Second, theories and models of visually guided action focus almost exclusively on optical invariants. This narrow view is incompatible with evidence from other studies indicating that observers sometimes use nonspecifying variables, and has hampered the development of a theory of perceptual learning in visually guided action.

5 A new framework for research on visually guided action

5.1 *Keeping the ideal state within the 'safe region'*

The aforementioned problem of action boundaries (see section 4) means that the observer's maximum action capabilities define critical boundaries separating possible from impossible actions. In fly-ball catching, for example, maximum running speed defines a critical boundary separating possible from impossible catches. When the ideal running speed is below the boundary, it can be kept below that boundary by running faster or slower. If the ideal running speed exceeds maximum running speed, then it will continue to increase even if maximum running speed is achieved, and the fielder will miss the ball. Thus, keeping the ideal state within the 'safe' region between zero and the maximum-possible state is both necessary and sufficient for successful performance. This implies an alternative to the error-nulling means of control suggested by existing models of visually guided action. Rather than making adjustments to null the difference between the current and ideal states, perhaps observers make adjustments to keep the ideal state within the safe region. When the ideal state is well within the safe region and gradually changing, it may not be necessary to make an adjustment, even if the current and ideal states are unequal. On the other hand, if the ideal state is close to the boundary of the safe region, an adjustment is necessary.

In fact, making adjustments to keep the ideal state within the safe region can actually give the appearance of error nulling. Figure 1a shows a sample trial from an experiment on visually guided braking reported in Fajen (2005). As the ideal deceleration (solid line) drifts towards the upper or lower boundary, the observer makes brake adjustments (dotted line) to reverse the direction of drift. The result is that the gap between the ideal and current deceleration shrinks. Thus, the behavior that researchers have interpreted as evidence of error nulling is arguably just a byproduct of keeping the ideal state within the safe region. In other words, the appearance of error nulling emerges automatically.

Unlike the traditional error-nulling framework in which there is a preferred state (eg the ideal deceleration, running speed, turning rate) in which the observer should strive to be at each moment, the new framework posits a safe region within which the observer should remain, but no preferred state within that safe region. One could argue that the new framework lacks a specific account (eg a control law) of how action is controlled within the safe region. But the fact is that there are an infinite number of possible trajectories through the safe region that correspond to successful performance. In braking, for example, one could adopt a conservative style and minimize collisions by making adjustments to keep the ideal deceleration near the bottom of the safe region, close to zero (see figure 1b). Or one could adopt a more aggressive style to minimize approach time by allowing the ideal deceleration to draw near the top of the safe region, close to maximum deceleration. Another possibility is to minimize boredom by allowing the ideal deceleration to oscillate erratically within the safe region. All of these are ways to perform the tasks successfully, so it is unlikely that there could be a single preferred state in which the observer should be across all

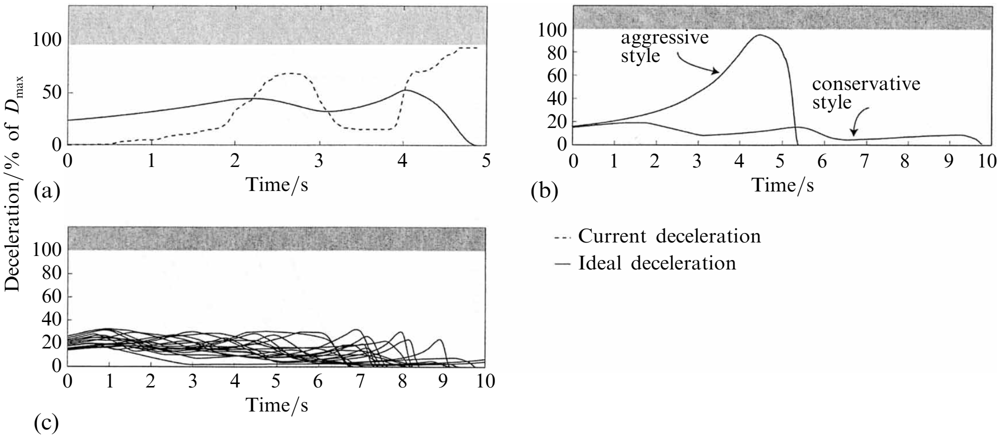

Figure 1. (a) Sample trial from an experiment on braking control showing current deceleration (dotted line) and ideal deceleration (solid line) as a function of time. The boundary separating the white and gray regions corresponds to the maximum deceleration of the brake. (b) Ideal deceleration as a function of time for two sample trials illustrating an aggressive and a conservative braking style. Ideal deceleration approaches maximum deceleration during the aggressive approach, increasing the risk of collision. On the conservative approach, ideal deceleration is kept well below maximum deceleration, but at the cost of a longer approach time. (c) Sample trials showing ideal deceleration constrained to the lower part of the safe region.

conditions and all individuals. The only constraint on successful performance that is common to all situations and all individuals is that the ideal state remains within the safe region. How action is actually controlled within the safe region depends on situation-specific factors. But all of these factors may be understood as additional 'softer' constraints on trajectories through the safe region. In the real-world, for example, the cost of collision and the inertial consequences of maintaining high deceleration rates are considerably more significant than they are in the laboratory. So during real-world driving, trajectories may be constrained to the lower part of the safe region (see figure 1c). The important point, however, is that these factors introduce additional constraints on trajectories; they do not prescribe a preferred trajectory. Adjustments are not made around an ideal state but rather within a safe region.

5.2 *The necessity of calibration*

To keep the ideal state within the safe region, one must perceive the ideal state relative to the maximum possible state. Information about the ideal state may be available in the optic array, but the maximum possible state is a property of the observer's body or vehicle for which there is no information. If these tasks are performed by keeping the ideal state within the safe region, and the critical boundary of the safe region defined by the limits of the observer's action capabilities is 'invisible', then how do observers know the location of the boundary well enough to control action? Clearly, some form of calibration must be involved.

Suppose a vehicle is approaching an obstacle 50 m ahead at a rate of 20 m s^{-1}. The ideal deceleration (ie the rate of deceleration that would bring the observer to a stop at the obstacle without making any further brake adjustments along the way) is equal to 4 m s^{-2}. Information in optic flow specifies ideal deceleration,[(5)] but, of course,

[(5)] In terms of spatial variables, ideal deceleration is equal to $\upsilon^2/2z$, where υ is speed and z is distance. Speed is optically specified by global optic flow rate (GOFR), which is the optical velocity of texture elements corresponding to the ground plane in a given visual direction (Larish and Flach 1990). υ/z is the inverse of time-to-contact, which is optically specified by $1/\tau$. Substituting GOFR for υ and $1/\tau$ for υ/z, the equation of ideal deceleration can be rewritten in terms of optical variables as GOFR$/2\tau$.

this does not mean that ideal deceleration is perceived in units of m s^{-2}. Meters and seconds are extrinsic units because they are defined independently of the observer, which makes them useful for researchers but completely arbitrary and meaningless to perceivers. For someone engaged in a visually guided action, the only units that are meaningful are those defined by one's action capabilities, which for the task of braking correspond to maximum deceleration. If the maximum deceleration of the vehicle in the example above is 10 m s^{-2} and the observer was well calibrated, then perceived ideal deceleration would be equal to 0.4, or 40% of maximum deceleration.

Perceiving ideal deceleration in units of maximum deceleration is tantamount to getting information and control in the same units. Just as ideal deceleration can be expressed as a percentage of maximum deceleration, brake displacement can be expressed as a percentage of maximum brake displacement. If the brake is designed so that the mapping from brake displacement to deceleration is proportional, then intrinsically scaled information about ideal deceleration specifies the range of brake positions necessary to keep the ideal deceleration within the safe region, without the need for 'knowledge' of controller dynamics. For example, if ideal deceleration is 40% of maximum deceleration and perceived as such by a calibrated observer, then the ideal deceleration could be kept within the safe region by applying 40% or more of maximum brake pressure.

Of course, maximum deceleration depends on many factors in addition to brake strength (eg load, surface friction, surface slope), some of which can change quickly. So practice and experience are necessary to recalibrate. From the present perspective, recalibration is achieved by adjusting the metric in which information is detected. The question that still remains is how observers discover the right metric. Although there is no information in optic flow about maximum action capabilities per se, the optical consequences of the observer's actions can be used to calibrate information to one's action capabilities. If the observer is properly calibrated, then the ideal deceleration (expressed as a percentage of maximum deceleration) will always drift away from brake position (expressed as a percentage of maximum brake position). This is illustrated in figure 1a. Ideal deceleration drifts upwards when brake position (or current deceleration) is less than ideal deceleration, and downwards when brake position is greater than ideal deceleration. If perceived ideal deceleration drifts toward brake position, then the observer is not properly calibrated. A rather unsettling example of this can occur the first time a driver tests out the brakes after hitching a heavy trailer. The driver may perceive that there is still time to stop (ie that ideal deceleration is less than 100% of maximum deceleration). But, if the additional load is large enough to significantly affect maximum deceleration, then perceived ideal deceleration may continue to increase even when maximum deceleration is applied. The upward drift in perceived ideal deceleration would indicate that the units in which ideal deceleration is perceived are too large (ie that the driver is calibrated to a stronger brake) and must be adjusted. Note that, if the brake is linear, the observer does not have to perform emergency stops to test out the maximum limits. The direction of drift in perceived ideal deceleration relative to brake position specifies whether or not the observer is calibrated regardless of brake position. It may, however, help to test out the maximum limits, because there is no uncertainty about brake position when the brake is in the maximum position. For positions in between zero and maximum, uncertainty about brake position could complicate attempts to calibrate.

Although the necessity of calibration was illustrated for the case of the braking task, the same principles apply to each of the other visually guided actions. Moreover, the hypothesis that information is detected in intrinsic units has some intriguing theoretical implications that reflect the way in which perception and action are deeply intertwined. When we say things like "I know my car's brake strength", the implicit assumption is that we have knowledge of the strength of the brake that is somehow represented and retrieved from memory when needed. This interpretation is consistent with a broader theoretical perspective according to which the guidance of action based on optic flow must be supplemented by various kinds of cognitive representations, such as knowledge of plant and controller dynamics (Loomis and Beall 1998). But, if observers calibrate to changes in brake strength by scaling information in intrinsic units, then the only sense in which an observer's knowledge of brake strength is represented is in terms of the metric in which information is detected. In other words, there is no representation of brake strength that exists independently of the detection of information—the representation is part of the mechanism that detects information. Likewise, if ideal deceleration is perceived in units of maximum deceleration, then changes in maximum deceleration will result in changes in perceived ideal deceleration (once recalibration has taken place). So the perception of ideal deceleration is not independent of one's capabilities to control braking. More broadly, for an observer engaged in a visually guided action, there is arguably no meaningful sense in which perceived 3-D space is represented independently of action, and no meaningful sense in which plant and controller dynamics are represented independently of the detection of information. An analogous argument has been made by Proffitt and colleagues (Proffitt et al 1995, 2003; Bhalla and Proffitt 1999), who showed that perceptual judgments of geographical slant are influenced by factors that affect one's behavioral potential (eg wearing a backpack).

5.3 *The necessity of perceptual learning*

In section 4, it was shown that approaching visually guided action as an error-nulling process thwarts attempts to investigate the possibility that people use nonspecifying variables, and flexibly attune to more effective optical variables with practice and experience. From the perspective that actions are controlled by keeping the ideal state within the safe region, the possible use of nonspecifying variables can be more easily approached.

Consider a scenario in which an observer is approaching an object at a constant speed. The ideal deceleration will increase until eventually it exceeds the maximum possible deceleration. Compare this scenario to another in which the observer approaches the same object at a faster constant speed. Again, the ideal deceleration will increase, but in this case it will exceed maximum deceleration sooner. One of the things that it means for an optical variable to specify ideal deceleration is that the value of such a variable when the ideal deceleration equals maximum deceleration is the same for both scenarios. More generally, if an optical variable specifies the ideal deceleration, then its value at the boundary is invariant over changes in such factors as approach speed and object size. This is illustrated in figure 2a, which shows the value of a specifying variable as a function of time for 25 simulated trials in which target size, initial distance, and initial speed all vary. The black dots indicate the point at which the ideal deceleration reached the maximum deceleration on each simulated trial. Note that the value of the optical variable at the critical boundary is the same, illustrating that the specifying variable is invariant across changes in size and speed at the critical boundary. In contrast, nonspecifying variables can take on different values at the boundary depending on approach speed and object size. Figure 2b shows the value of expansion rate (ie a nonspecifying variable) as a function of time for the same 25 simulated trials used in figure 2a. Again, the dots indicate the point at which the ideal deceleration reached the maximum deceleration. Unlike the specifying variable in figure 2a, the value of the nonspecifying variable at the critical boundary changes from trial to trial as size and speed change, indicating that expansion rate is not invariant over size and speed.

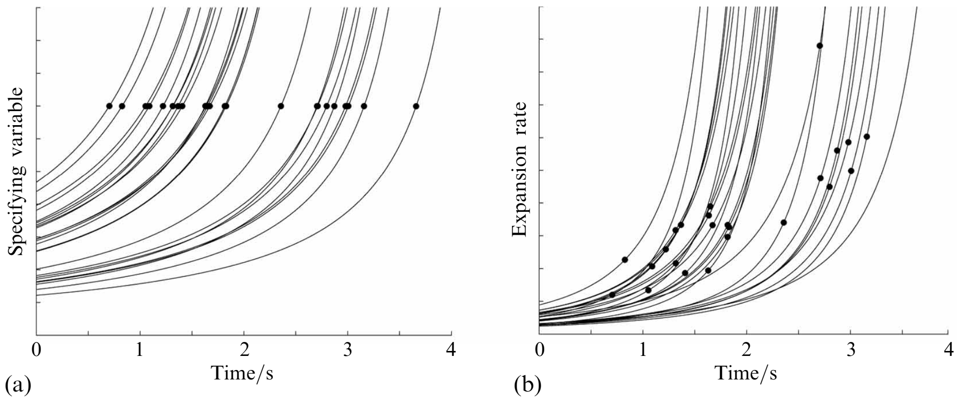

Figure 2. Plots of (a) a specifying variable; and (b) a nonspecifying variable, as a function of time for 25 simulated trials using different initial conditions. Initial distance varied between 40 and 50 m, speed varied between 10 and 20 m s^{-1}, and target radius varied between 0.25 and 0.75 m. The black dots indicate the point in each simulated trial at which the ideal deceleration exceeded the maximum deceleration of the brake (7.0 m s^{-2}).

Because nonspecifying variables are not invariant over size and speed, performance based on such variables will be systematically biased. To visualize the pattern of bias, it is useful to plot the value of the chosen optical variable at the moment that the ideal and maximum decelerations are equal as a function of size and speed (see Appendix for the step-by-step procedure used to create these plots). The specifying variable does not vary across size and speed at the action boundary, and so it forms a flat surface perpendicular to the vertical axis (see figure 3a). In contrast, expansion rate does vary across size and speed at the action boundary, and hence it forms a curved surface (see figure 3b). At the moment that the ideal and maximum decelerations are equal,

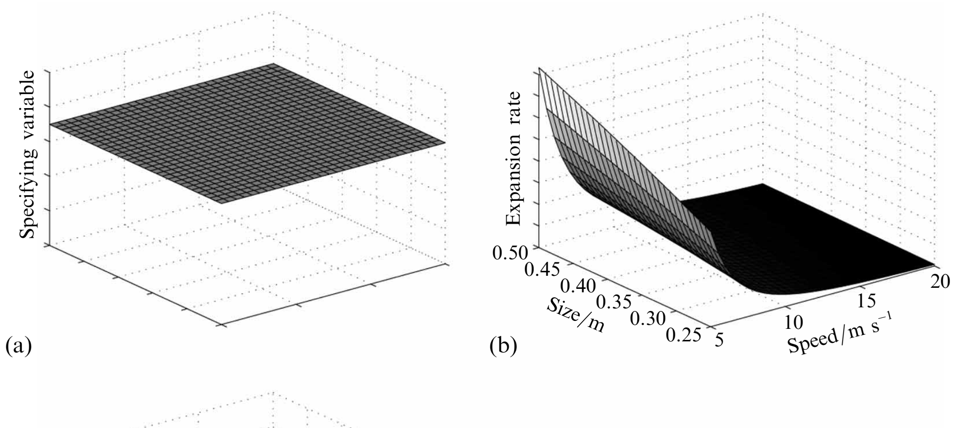

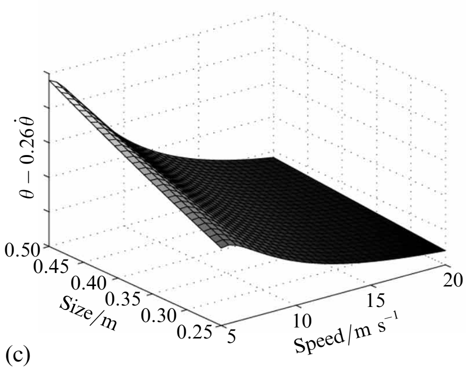

Figure 3. Surface plots showing (a) the value of a specifying variable; (b) an unreliable nonspecifying variable; and (c) a reliable nonspecifying variable at the critical boundary as a function of speed and target size.

expansion rate is lower for smaller sizes and for faster speeds. If braking is controlled by keeping the perceived ideal deceleration between zero and 1.0 (in intrinsic units to which the observer is calibrated), and ideal deceleration is perceived on the basis of $\dot{\theta}$, then performance will be biased owing to a tendency to increase brake pressure later or more gradually when size is small and when speed is fast. To avoid collisions when such conditions are encountered, the scaled expansion rate should be less than 1.0 at the action boundary even when size is small and speed is fast, which means that expansion rate must be scaled in relatively small units. While this is adaptive for smaller sizes and faster speeds, it will result in the opposite bias for larger sizes and slower speeds. That is, because expansion rate is considerably higher in this region of the space, observers should be more likely to apply brake pressure unnecessarily. Thus, the use of expansion rate may result in a safer, conservative style of braking, but may also lead to unnecessarily slow approaches.

Of course, the degree of bias depends on the range of conditions encountered. When speed and size do not vary much, the range of values that expansion rate can take on at the action boundary is more limited. Furthermore, there are other nonspecifying variables that vary more or less than expansion rate. Figure 3c shows the surface corresponding to a linear combination of angle and expansion rate. Compared to expansion rate, the surface in figure 3c is relatively flat throughout most of the range. The upward slope of the surface indicates that relying on $a\dot{\theta} + b\theta$ may lead to unnecessary increases in brake pressure at slower speeds. But the slope is not nearly as steep as the slope in figure 3b. Thus, $a\dot{\theta} + b\theta$ is more effective than $\dot{\theta}$ alone, but not as effective as the specifying variable. Other effective nonspecifying variables are likely to be discovered by exploring their variability at the critical boundary.

The surface plots in figure 3 show how empirically testable predictions about the use of nonspecifying variables can more easily be made once the error-nulling assumption is rejected and replaced with the assumption that actions are controlled by keeping the perceived ideal state within the safe region. This opens the door for the investigation of nonspecifying variables and attunement to more effective variables with practice in visually guided action.

To summarize, the widespread view that visually guided action is an error-nulling process has diverted attention away from two critical aspects of visual control: (i) successful performance requires observers to be extremely sensitive to the limitations of their action capabilities; and (ii) observers sometimes rely on nonspecifying variables, and the ability to adapt to changing conditions and to improve performance reflects a process of perceptual attunement (or convergence) onto more effective optical variables. In this section, it was shown how both issues could be addressed by abandoning the error-nulling and single-optical-invariant assumptions, and replacing them with the principles of scaling information to action (ie calibration), and flexible attunement. Within this framework, successful visually guided action is a process by which well-calibrated observers keep the ideal state within the safe region. The ability to adapt to changing conditions and improve performance with practice reflects a process of perceptual learning, which involves attunement to more effective optical variables.

6 Affordances and visually guided action

I began this article by claiming that successful action depends on the ability to perceive the behavioral possibilities of one's environment; that is, in both routine and skilled tasks, one must perceive what actions are possible and what actions are not possible. This is not a novel claim. Rather, it follows from one of the core assumptions of Gibson's (1986) ecological approach, that the environment is perceived in terms of what the observer can and cannot do within it; that is, to see something is to see what to do with it. Gibson coined the term 'affordance' to provide a description of the environment in terms of its possibilities for action. Whether or not an action is possible depends as much on the observer's body dimensions and action capabilities as it does on the properties of the environment. So an affordance is an invariant combination of properties of substances and surfaces, but must also be taken with respect to the animal. For example, a 1 m ledge affords stepping down for an adult but not for a small child, and a wall affords climbing for many insects but not for humans. That the environment affords such actions is, in and of itself, a trivial statement. The innovation in Gibson's theory of affordances is that these are real properties of the animal–environment system that are directly perceived (Turvey 1992). From an ecological perspective, affordances are what the environment of an animal means to that animal. In other words, meaning does not originate in the observer's head, but is out there to be perceived by detecting information about affordances.

Because the theory of affordances was developed by Gibson and others to account for the inseparability of perception and action, one would expect there to be a consistency between the existing framework for research on visually guided action and the view that successful action depends on the perception of affordances. Minimally, one would expect the affordance for each visually guided action to be clearly identified in the literature, and the relevant optical variable in each model to specify the affordance. Surprisingly, this is not the case. Affordances are rarely mentioned in papers on visually guided action, and the properties specified by optical variables in models of visually guided actions are not possibilities for action! Recall that a common assumption of existing models of visually guided action is that observers must monitor the sufficiency of their current state. In other words, outfielders must perceive the property "I am (or am not) running fast enough to catch the ball" and cyclists must perceive the property "I am (or am not) turning enough to pass through that opening". These are the properties that are specified by the optical variables used in models of fly-ball catching and steering, but they are not affordances. The affordance for fly-ball catching is "It is (or is not) possible to catch the ball". This is not the same as "I am (or am not) running fast enough to catch the ball" and is not specified by the optical acceleration of the ball [ie the optical variable in Chapman's (1968) OAC model of fly-ball catching]. Likewise, the affordance for steering is "It is (or is not) possible to reach that goal", which is not the property that is specified by optical variables used in models of steering. Thus, if the property that cyclists perceive is "I am (or am not) turning enough to reach that goal", then they do not perceive the affordance, and there is a fundamental inconsistency between models of visually guided action and the theory of affordances.

This discrepancy can once again be resolved by abandoning the assumption that people make adjustments to null the error between the current and ideal state, and adopting the alternative framework outlined above. In each of the tasks discussed throughout this paper, an action is possible as long as the ideal state does not exceed the observer's maximum action capabilities. A fly ball affords catching as long as the ideal running speed does not exceed the fielder's maximum running speed. An obstacle affords avoiding as long as the ideal deceleration does not exceed the maximum deceleration. Thus, when one detects information about the ideal state relative to (or in units of) maximum action capabilities, one perceives whether or not an action is possible—that is, one perceives the affordance.

Resolving the discrepancy between visually guided action and affordances opens the door for the investigation of a class of affordances that has so far received very little attention. Most affordance research focuses on body-scaled affordances—possibilities for action that are defined by one's body dimensions—such as stair climbing (Warren 1984), sitting (Mark 1987), reaching (Carello et al 1989), and passing through an aperture (Warren and Whang 1987). These affordances are body-scaled because they are defined by dimensions of the environment taken with respect to dimensions of the body, such as leg length, arm length, and shoulder width. In contrast, the affordances in visually guided actions are defined by one's action capabilities. Whether or not a ball is catachable depends not on the fielder's leg length or eye height (although it is likely to be correlated with these dimensions), but on the fielder's maximum running speed (Oudejans et al 1996). Whether or not a goal is reachable depends on the observer's maximum turning rate. Just as affordances for stair climbing and sitting must be scaled to the actor's body dimensions, so must the affordances for catching and braking be scaled to the actor's action capabilities. In section 7, I discuss the research implications of espousing this approach.

6.1 *Smart perceptual instruments*
If one adopts the ecological stance that successful action depends on the perception of affordances, then the primary component of our perceptual experience must be whether or not an action is possible. When braking, for example, our most immediate awareness is that it is or is not still possible to stop. Our awareness of spatial properties of the environment, such as distance, speed, size, time-to-contact, and deceleration is arguably irrelevant to the control of braking, and perceiving them may not be a necessary prerequisite to perceiving whether or not it is possible to stop. This may seem counterintuitive, because the ideal deceleration is a higher-order property that can be derived from a combination of these lower-order properties (eg $v^2/2z$). But the assumption that the perception of higher-order properties necessarily depends on the perception of lower-order properties is invalid. Runeson (1977) used the example of a polar planimeter to illustrate how 'smart' instruments can be designed to measure higher-order properties without first measuring lower-order properties. A polar planimeter is a device that was invented in 1854 by Jacob Amsler to measure the area of an irregularly shaped region (eg a piece of land on a map). The interesting thing about the polar planimeter is that it directly measures the area of such regions without taking any measurements of the length and width of the region and without performing any calculations or making any inferences. The point of this argument is that if instruments can be designed to measure higher-order properties (eg area) without first measuring lower-order properties (eg length and width), then perhaps perceptual systems can organize themselves into devices for measuring higher-order properties of the environment. Specifically, a properly tuned, well-calibrated perceptual system could measure the ideal deceleration (in units of maximum deceleration) without first measuring any of the lower-order properties such as distance, speed, or time-to-contact.

7 Implications for designing experiments and models of visually guided action

The acknowledgment that calibration, perceptual learning, and the perception of affordances play essential roles has broad implications for the goals and objectives of research on visually guided action. The purpose of this section is to make these implications explicit, and to illustrate how the proposed framework may influence the design of experiments, and the analysis and modeling of data.

(i) A primary goal of most of the previous research on visually guided action described in section 3 is to identify the optical variables that specify the sufficiency of one's current state. For example, $\dot{\tau}$ specifies the sufficiency of one's current deceleration to avoid a collision, the change in bearing angle specifies the sufficiency of one's current walking speed (or direction) to intercept a moving target, the optical acceleration of the elevation angle of a fly ball specifies the sufficiency of the outfielder's current running speed, and the locomotor flow line specifies sufficiency of one's current turning rate to steer to a stationary goal. Throughout this paper, it has been suggested that observers do not rely upon information about the sufficiency of their current state, but rather upon information about the ideal state (relative to one's maximum action capabilities). Thus, one implication is that researchers should look for optical variables that specify the ideal state, and design experiments to test whether observers use these variables. For example, $2\sin\theta/z$ specifies the path curvature required to steer to a stationary goal (Fajen 2001; Wann and Land 2000), and $\mathrm{GOFR}/2\tau$ specifies the deceleration required to avoid a collision with an obstacle in the path of motion (Yilmaz and Warren 1995). In addition, researchers must also consider optical variables that closely correlate with, but do not specify, the ideal state across the range of conditions used in their experiments. The method described in section 5 of this paper can be used to estimate the effectiveness of nonspecifying variables and to make empirically testable predictions.

(ii) A second important goal for future research is to look for evidence that observers are sensitive to the boundaries of their action capabilities. When the ideal state is well within one's maximum action capabilities, the consequences of each adjustment are relatively insignificant, and one would expect a great deal of movement variability—observers may make small adjustments in either direction, or no adjustments at all. As the ideal state approaches the boundary, an adjustment must be made. If observers are sensitive to the boundaries of their action capabilities, then one would expect the variability of adjustments to decrease near the boundary. Fajen (2005) demonstrated this in a braking task by showing that, when the ideal deceleration was below the maximum deceleration, observers were just as likely to increase brake pressure as they were to decrease or hold brake pressure. Near the boundary where the ideal deceleration was equal to the maximum deceleration, observers almost always increased brake pressure.

(iii) Researchers should look for evidence that information is scaled in units of the observer's maximum action capabilities—that is, for evidence of calibration and recalibration. One effective technique for testing calibration is to manipulate the observer's action capabilities. For example, in a virtual or simulated environment, researchers could manipulate the observer's maximum running speed, turning rate, or deceleration by changing the visual gain. If information is scaled in units of the observer's maximum action capabilities, then performance should be comparable when the relevant action measures are plotted as a function of the ideal state expressed in units of the maximum action capabilities. For example, I (Fajen, in press) recently found that observers calibrated to different brake strengths initiate brake adjustments at the same mean ideal deceleration when ideal deceleration is expressed as a percentage of maximum deceleration for each brake strength condition.

(iv) The proposed approach provides a convenient framework within which to investigate the learning of a visually guided action as a process by which observers detect more effective optical variables with practice. Researchers should design experiments so that the performance of each subject can be analyzed at various stages of learning. During the early stages, one would expect there to be large individual differences in the use of optical variables, and many observers may use nonspecifying variables. Possible variables can be eliminated by looking for systematic biases, and for evidence that observers are insensitive to the critical boundary across a wide range of conditions. With practice, observers may converge on similar, more effective variables. As observers learn to detect a specifying variable, systematic biases should diminish.

(v) Another important goal is to look for evidence that observers adapt to the range of conditions that they encounter by learning to detect different optical variables. In constrained environments, the effectiveness of some lower-order nonspecifying variables may be indistinguishable from the effectiveness of specifying variables. For example, recall that, when the size of a target does not vary, the effectiveness of a linear combination of optical angle and expansion rate for controlling braking is comparable to the effectiveness of the specifying variable. Observers who practice in environments with fixed-size targets may not learn to use a specifying variable because feedback about performance with the linear combination is adequate. As conditions change, observers may adapt by learning to detect more effective optical variables. Researchers should test for such adaptation by designing experiments in which performance is compared across different ranges of practice conditions.

(vi) Researchers may also explore individual differences in style of control within the proposed framework. Keeping the ideal state well below the critical boundary is the sign of a safe, conservative style of control. In braking, for example, a cautious driver would keep the ideal deceleration well below the maximum deceleration, but at the cost of slower, more time consuming approaches. A more aggressive style of control would be signified by allowing the ideal state to approach the critical-action boundary. In braking, this would lead to quicker approaches, but at the cost of greater risk. If the ideal deceleration accidentally exceeds maximum deceleration, then a collision is unavoidable. Highly skilled individuals who are extremely sensitive to the location of the boundary may be capable of keeping the ideal state near the boundary without allowing it to escape.

(vii) The necessity of calibration means that researchers need to be cautious about interpreting data from experiments that require participants to make passive judgments about tasks that normally require active control. For example, Andersen et al (1999) asked participants to watch displays simulating constant-deceleration approaches to a stop sign. The displays terminated before the stop sign was reached, and participants were asked to judge whether they would have collided or stopped short. Andersen et al found that judgments were influenced by approach speed, sign size, and edge rate, independently of the value of $\dot{\tau}$ when the displays terminated. One problem with this task is that observers were asked to make judgments about the sufficiency of their current deceleration, which I argue is not the relevant property for controlling braking. If braking is controlled by detecting information about the ideal deceleration, then perhaps observers should have been asked to judge whether or not it was still possible to stop when the displays were terminated. To make such a judgment, however, information about the ideal deceleration must be scaled in units of the maximum deceleration. Feedback about the correctness of the participant's 'stop-ability' judgment could be provided after each trial, but most of the information for calibration during actual braking is likely to come from seeing the optical consequences of individual brake adjustments. Thus, the second problem with passive judgment tasks is that the normally available feedback that is necessary to scale information into action-relevant units is unavailable.

(viii) Lastly, new models of visually guided action will be required which include a necessary role for calibration and perceptual learning. Table 1 provides an outline for models of each of the visually guided actions that were reviewed in section 3. The relevant control variable whose ideal state must be perceived is listed in the second column for each task. For example, running speed is the relevant control variable for catching a fly ball because the outfielder must perceive the running speed required to arrive at the landing location at the right time. Optical variables that specify the ideal state of the control variable are listed in the third column for tasks in which such variables have already been derived. Nonspecifying variables should also be considered. The lower and upper boundaries of the safe region for the ideal state of the relevant control variable are defined in columns four and five, respectively. The upper boundary is always defined by the maximum action capability, and the lower boundary is always zero.

Table 1. Variables for revised models of visually guided action.

Task	Control variable	Optical variable	Lower boundary	Upper boundary
Steering to a stationary goal	path curvature	$2\sin\theta/z$	zero curvature	maximum path curvature
Intercepting on foot	running speed	?	zero running speed	maximum running speed
Catching a fly ball	running speed	?	zero running speed	maximum running speed
Intercepting by hand	hand speed	?	zero hand speed	maximum hand speed
Braking	deceleration	$\mathrm{GOFR}/2\tau$	zero deceleration	maximum deceleration

8 Conclusion

The primary purpose of this paper is to offer a new framework for research on visually guided action, with the hope that it will initiate a second round of thinking about tasks such as steering, intercepting, fly-ball catching, and braking. Current research on visually guided action is shaped by the assumption that observers rely on a single optical invariant for each task that specifies the sufficiency of their current state, and make adjustments to null the error between the current state and the state required to successfully perform the task without making any further adjustments along the way. One might argue that a considerable amount of important research has been conducted within this framework, and that a new framework is not needed. Indeed, we have learned a great deal about fundamental tasks such as steering, braking, and intercepting over the past twenty-five years by approaching visually guided action as an error-nulling process. The existing framework has reshaped our understanding of visual space perception, inspired researchers to look for higher-order properties of the optic-flow field, and arguably delivered on its promise to explain the performance of these complex tasks without the appeal to an internal model of the environment. Nonetheless, the assumptions that underlie the existing framework were chosen to make these aspects of visually guided action scientifically approachable. At the same time, they have served as conceptual barriers to understanding essential aspects of visually guided action. It is only by rejecting the assumption that visual guidance is an error-nulling process that one can begin to explore the necessary role of calibration, perceptual learning, and perceiving possibilities for action. Within the alternative framework proposed in this paper, one must learn to detect and calibrate optical information about the boundary separating possible from impossible actions. Thus, successful performance in visually guided action is a matter of perceiving possibilities for action, and behaving so as to keep the desired action within the range of possible actions.

Acknowledgment. This research was supported by the National Science Foundation (BCS 0236734).

References

Andersen G J, Cisneros J, Atchley P, Saidpour A, 1999 "Speed, size, and edge-rate information for the detection of collision events" *Journal of Experimental Psychology: Human Perception and Performance* **25** 256 – 269

Bardy B G, Warren W H, 1997 "Visual control of braking in goal-directed action and sport" *Journal of Sports Sciences* **15** 607 – 620

Beall A C, Loomis J M, 1996 "Visual control of steering without course information" *Perception* **25** 481 – 494

Bernstein N, 1967 *The Coordination and Regulation of Movement* (London: Pergamon Press)

Bhalla M, Proffitt D R, 1999 "Visual – motor recalibration in geographical slant perception" *Journal of Experimental Psychology: Human Perception and Performance* **25** 1076 – 1096

Bootsma R J, Fayt V, Zaal F T, Laurent M, 1997 "On the information-based regulation of movement: What Wann (1996) may want to consider" *Journal of Experimental Psychology: Human Perception and Performance* **23** 1282 – 1289
Caird J K, Hancock P A, 1994 "The perception of arrival time for different oncoming vehicles at an intersection" *Ecological Psychology* **6** 83 – 109
Carello C, Grosofsky A, Reichel F D, Solomon H Y, Turvey M T, 1989 "Visually perceiving what is reachable" *Ecological Psychology* **1** 27 – 54
Chapman S, 1968 "Catching a baseball" *American Journal of Physics* **36** 368 – 370
DeLucia P R, 1991 "Pictorial and motion-based information for depth perception" *Journal of Experimental Psychology: Human Perception and Performance* **17** 738 – 748
DeLucia P R, Warren R, 1994 "Pictorial and motion-based depth information during active control of self-motion: size-arrival effects on collision avoidance" *Journal of Experimental Psychology: Human Perception and Performance* **20** 783 – 798
Duchon A P, Warren W H Jr, 2002 "A visual equalization strategy for locomotor control: of honeybees, robots, and humans" *Psychological Science* **13** 272 – 278
Duchon A P, Warren W H, Kaelbling L P, 1998 "Ecological robotics" *Adaptive Behavior* **6** 473 – 507
Fajen B R, 2001 "Steering toward a goal by equalizing taus" *Journal of Experimental Psychology: Human Perception and Performance* **27** 953 – 968
Fajen B R, 2005 "Calibration, information, and control strategies for braking to avoid a collision" *Journal of Experimental Psychology: Human Perception and Performance* **31** in press
Fajen B R, in press, "The scaling of information to action in visually guided braking" *Journal of Experimental Psychology: Human Perception and Performance*
Fajen B R, Warren W H, 2003 "Behavioral dynamics of steering, obstacle avoidance, and route selection" *Journal of Experimental Psychology: Human Perception and Performance* **29** 343 – 362
Fajen B R, Warren W H, 2004 "Visual guidance of intercepting a moving target on foot" *Perception* **33** 689 – 715
Fajen B R, Warren W H, Temizer S, Kaelbling L P, 2003 "A dynamical model of visually-guided steering, obstacle avoidance, and route selection" *International Journal of Computer Vision* **54** 13 – 34
Flach J M, Warren R, Garness S A, Kelly L, Stanard T, 1997 "Perception and control of altitude: splay and depression angles" *Journal of Experimental Psychology: Human Perception and Performance* **23** 1764 – 1782
Gibson J J, 1958 "Visually controlled locomotion and visual orientation in animals" *British Journal of Psychology* **49** 182 – 194
Gibson J J, 1966 *The Senses Considered as Perceptual Systems* (Boston, MA: Houghton Mifflin)
Gibson J J, 1986 *The Ecological Approach to Visual Perception* (Hillsdale, NJ: Lawrence Erlbaum Associates)
Gilden D L, Proffitt D R, 1989 "Understanding collision dynamics" *Journal of Experimental Psychology: Human Perception and Performance* **15** 372 – 383
Gilden D L, Proffitt D R, 1994 "Heuristic judgment of mass ratio in two-body collisions" *Perception & Psychophysics* **56** 708 – 720
Holst E von, 1954 "Relations between the central nervous system and the peripheral organs" *British Journal of Animal Behaviour* **2** 89 – 94
Jacobs D M, Michaels C F, 2002 "On the apparent paradox of learning and realism" *Ecological Psychology* **14** 127 – 139
Jacobs D M, Runeson S, Michaels C F, 2001 "Learning to visually perceive the relative mass of colliding balls in globally and locally constrained task ecologies" *Journal of Experimental Psychology: Human Perception and Performance* **27** 1019 – 1038
Kamp J van der, Savelsbergh G J, Smeets J, 1997 "Multiple information sources in interceptive timing" *Human Movement Science* **16** 787 – 821
Kim N G, Turvey M T, 1999 "Eye movements and a rule for perceiving direction of heading" *Ecological Psychology* **11** 223 – 248
Larish J F, Flach J M, 1990 "Sources of optical information useful for perception of speed of rectilinear self-motion" *Journal of Experimental Psychology: Human Perception and Performance* **16** 295 – 302
Lee D N, 1976 "A theory of visual control of braking based on information about time-to-collision" *Perception* **5** 437 – 459
Lee D N, Lishman R, 1977 "Visual control of locomotion" *Scandinavian Journal of Psychology* **18** 224 – 230

Lenoir M, Musch E, Janssens M, Thiery E, Uyttenhove J, 1999 "Intercepting moving objects during self-motion" *Journal of Motor Behavior* **31** 55 – 67
Loomis J M, Beall A C, 1998 "Visually controlled locomotion: Its dependence on optic flow, three-dimensional space perception, and cognition" *Ecological Psychology* **10** 271 – 285
McBeath M K, Shaffer D M, Kaiser M K, 1995 "How baseball outfielders determine where to run to catch fly balls" *Science* **268** 569 – 573
McLeod P, Reed N, Dienes Z, 2001 "What we do not know about how people run to catch a ball" *Journal of Experimental Psychology: Human Perception and Performance* **27** 1347 – 1355
Mark L S, 1987 "Eyeheight-scaled information about affordances: A study of sitting and stair climbing" *Journal of Experimental Psychology: Human Perception and Performance* **13** 361 – 370
Michaels C F, Carello C, 1981 *Direct Perception* (Englewood Cliffs, NJ: Prentice-Hall)
Michaels C F, Vries M M de, 1998 "Higher order and lower order variables in the visual perception of relative pulling force" *Journal of Experimental Psychology: Human Perception and Performance* **24** 526 – 546
Michaels C F, Oudejans R R, 1992 "The optics and actions of catching fly balls: Zeroing out optical acceleration" *Ecological Psychology* **4** 199 – 222
Michaels C F, Zeinstra E B, Oudejans R R, 2001 "Information and action in punching a falling ball" *Quarterly Journal of Experimental Psychology A* **54** 69 – 93
Montagne G, Laurent M, Durey A, Bootsma R, 1999 "Movement reversals in ball catching" *Experimental Brain Research* **129** 87 – 92
Oudejans R R, Michaels C F, Bakker F C, Dolne M A, 1996 "The relevance of action in perceiving affordances: perception of catchableness of fly balls" *Journal of Experimental Psychology: Human Perception and Performance* **22** 879 – 891
Peper L, Bootsma R J, Mestre D R, Bakker F C, 1994 "Catching balls: how to get the hand to the right place at the right time" *Journal of Experimental Psychology: Human Perception and Performance* **20** 591 – 612
Proffitt D R, Bhalla M, Grossweiler R, Midgett J, 1995 "Perceiving geographical slant" *Psychonomic Bulletin & Review* **2** 409 – 428
Proffitt D R, Stefanucci J, Banton T, Epstein W, 2003 "The role of effort in perceiving distance" *Psychological Science* **14** 106 – 112
Runeson S, 1977 "On the possibility of 'smart' perceptual mechanisms" *Scandinavian Journal of Psychology* **18** 178 – 179
Runeson S, Vedeler D, 1993 "The indispensability of precollision kinematics in the visual perception of relative mass" *Perception & Psychophysics* **53** 617 – 632
Rushton S K, Harris J M, Lloyd M, Wann J P, 1998 "Guidance of locomotion on foot uses perceived target location rather than optic flow" *Current Biology* **8** 1191 – 1194
Shaffer D M, McBeath M K, 2002 "Baseball outfielders maintain a linear optical trajectory when tracking uncatchable fly balls" *Journal of Experimental Psychology: Human Perception and Performance* **28** 335 – 348
Smith M R, Flach J M, Dittman S M, Stanard T, 2001 "Monocular optical constraints on collision control" *Journal of Experimental Psychology: Human Perception and Performance* **27** 395 – 410
Srinivasan M V, 1998 "Insects as Gibsonian animals" *Ecological Psychology* **10** 251 – 270
Tresilian J R, 1999 "Visually timed action: Time-out for 'tau'?" *Trends in Cognitive Sciences* **3** 301 – 309
Turvey M T, 1992 "Affordances and prospective control: An outline of the ontology" *Ecological Psychology* **4** 173 – 187
Turvey M T, Fitch H L, Tuller B, 1982 "The Bernstein perspective: I. The problems of degrees of freedom and context-conditioned variability", in *Human Motor Behavior: An Introduction* Ed. J A S Kelso (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 239 – 252
Turvey M T, Shaw R E, Reed E S, Mace W M, 1981 "Ecological laws of perceiving and acting: In reply to Fodor and Pylyshyn (1981)" *Cognition* **9** 139 – 195
Wann J P, Land M, 2000 "Steering with or without the flow: is the retrieval of heading necessary?" *Trends in Cognitive Sciences* **4** 319 – 324
Wann J P, Swapp D K, 2000 "Why you should look where you are going" *Nature Neuroscience* **3** 647 – 648
Warren W H, 1984 "Perceiving affordances: Visual guidance of stair climbing" *Journal of Experimental Psychology: Human Perception and Performance* **10** 683 – 703
Warren W H, 1988 "Action modes and laws of control for the visual guidance of action", in *Movement Behavior: The Motor – Action Controversy* Eds O G Meijer, K Roth (Amsterdam: North-Holland) pp 339 – 379

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

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

Warren W H, Mestre D R, Blackwell A W, Morris M W, 1991 "Perception of circular heading from optical flow" *Journal of Experimental Psychology: Human Perception and Performance* **17** 28 – 43

Warren W H, 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

Yilmaz E H, Warren W H, 1995 "Visual control of braking: a test of the tau hypothesis" *Journal of Experimental Psychology: Human Perception and Performance* **21** 996 – 1014

Zaal F T, Michaels C F, 2003 "The information for catching fly balls: judging and intercepting virtual balls in a CAVE" *Journal of Experimental Psychology: Human Perception and Performance* **29** 537 – 555

Appendix

This section describes the step-by-step procedure used to create the surface plots in figure 3. The method is illustrated for three variables used in the context of the braking task. Variable 1 is $\mathrm{GOFR}/2\tau$, where GOFR is the 'global optic flow rate' and τ is the optical variable 'tau'. Global optic flow rate is the optical velocity of ground surface texture elements in a given visual direction, and is proportional to observer speed assuming constant eye height (Larish and Flach 1990). Thus $\mathrm{GOFR}/2\tau$ is proportional to the ideal deceleration ($v^2/2z$), and can be considered a specifying variable. Variable 2 is $\dot{\theta}$, or the rate of optical expansion of the target, and will be used as an example of an unreliable nonspecifying variable. Variable 3 is $a\dot{\theta} + b\theta$, which is a weighted combination of optical angle and expansion rate, and will be used as an example of a reliable nonspecifying variable.

(i) The first step is to select an optical variable and define it in terms of the relevant spatial variables [distance (z), speed (v), and sign size (r)].

Variable 1: $\mathrm{GOFR}/2\tau = kv^2/2z$ (where k is a constant that scales v to GOFR)

Variable 2: $\dot{\theta} = \dfrac{-2rv}{r^2 + z^2}$

Variable 3: $a\dot{\theta} + b\theta = a\left(\dfrac{-2rv}{r^2 + z^2}\right) + b\left(2\arctan\dfrac{r}{z}\right)$

(ii) Next, specify the critical boundary in terms of spatial variables.

$$d_{\mathrm{ideal}} = \frac{v^2}{2z} = D_{\max} \text{ (where } D_{\max} \text{ is the maximum deceleration)}$$

(iii) Substitute the critical boundary equation from step (ii) into the equations from step (i). The resulting equation describes the possible values that each optical variable can take on at the boundary.

Variable 1: $\mathrm{GOFR}/2\tau = kD_{\max}$

Variable 2: $\dot{\theta} = -2rv \Bigg/ \left[r^2 + \left(\dfrac{v^2}{2D_{\max}}\right)^2\right]$

Variable 3: $a\dot{\theta} + b\theta = a\left\{-2rv \Bigg/ \left[r^2 + \left(\dfrac{v^2}{2D_{\max}}\right)^2\right]\right\} + b\left[2\arctan\left(r \Bigg/ \dfrac{v^2}{2D_{\max}}\right)\right]$

(iv) Stipulate the range of conditions that the observer will encounter.

$0 < z < 50$ m

$0 < v < 20$ m s^{-1}

$0.25 < r < 0.50$ m

Plot the variable as a function of its components across the range specified in step (iv) (see figure 3).

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION
VOLUME 34 2005

www.perceptionweb.com