Ecological Robotics: Controlling Behavior with Optical Flow

Andrew P. Duchon
Department of Cognitive and Linguistic Sciences
Brown University
Providence, RI 02912
duchon@cog.brown.edu

William H. Warren
Department of Cognitive and Linguistic Sciences
Brown University
Providence, RI 02912
Bill_Warren@Brown.edu

Leslie Pack Kaelbling
Department of Computer Science
Brown University
Providence, RI 02912
lpk@cs.brown.edu

Abstract

There are striking parallels between ecological psychology and new trends in robotics and computer vision, particularly regarding how agents interact with the environment. We present some ideas from ecological psychology, including control laws using optical flow, affordances and action modes, and describe our implementation of these concepts in a small mobile robot which can avoid obstacles and play tag solely using optical flow. This work ties in with those of others arguing for a methodological approach in robotics which foregoes a central model/planner. Ecological psychology may not only contribute to robotics, but robotic implementations in turn provide a test bed for ecological principles and sources of ideas which could be tested in animals and humans.

Introduction

Classical symbolic systems have proven to be powerful in modeling some aspects of human cognition. However, the things humans do easily (e.g., recognizing patterns, moving around in the world, speaking) are difficult to explain and mimic with a symbolic approach. (We take the symbolic approach to mean symbolic/syntactic processing only, despite arguments for its all-inclusiveness (Vera & Simon, 1993).) Within both the cognitive science and AI communities an increasing dissatisfaction with this approach has led to the emergence of alternative approaches. In cognitive science, ecological psychology and especially connectionism are influential, and in the fields of robotics and computer vision, non-symbolic and situated architectures are having a strong impact.

In this paper we explore the similarities between ecological psychology and these new trends in robotics and computer vision. For its own reasons, ecological psychology has avoided explanations of perception and action which require a central model/planner and has promoted tighter binding between perception and action through the concepts of control laws, affordances, and action modes. We present these ideas and relate some of our work implementing them in a small mobile robot.

Behavior-Based Robotics and Active Vision

In the past few years, a new approach has developed in robotics called *behavior-based robotics* (Brooks, 1991a). Although earlier work on mobile robots had some successes, e.g., SRI's Shakey (Nilsson, 1984) and Moravec's (1981) CART, they were generally of the "sense-model-plan-act" variety, requiring intense computation for inferring the location and identity of objects, updating a central world model, and planning a course of action to achieve some defined goal state. In contrast, the new approach (Brooks, 1991b) attempts to build up robots through networks of simple, fully functional behaviors mapping sensors to actuators, with no central model. Complex behavior emerges from the dynamic interaction between the agent with its simple mappings and the environment, producing what appears to be goal-directed action.

Most of the work in robotics uses sensors other than vision. Sonar, infra-red detectors, laser-light stripers, and dead-reckoning provide metric distance information and traditionally, the robot uses this information to place itself at a particular point in its world model and to plan a metric path through the environment. The typical role of vision in these robots is to create or augment the model. Using visual information as simply another way to obtain metric values allows one to treat computer vision as a separate task, one of scene analysis: creating a description of the three-dimensional world from two-dimensional images. The numerous means of constructing such models (e.g., shape from shading, structure from motion), as they are presently formulated, are often ill-posed problems requiring assumptions and noise models which do not generalize to real-world vision (Aloimonos & Rosenfeld, 1991). However, with active control of the visual system (*active vision*) these problems become well-posed, usually with unique solutions and a few reasonable assumptions (Aloimonos & Rosenfeld, 1991; Ballard, 1991). *Purposive* or *animate vision* (Aloimonos, 1993; Ballard & Brown, 1992) goes one step further than finding better solutions to the old problems; rather it poses the question, "What is vision for?" (Ballard, 1991). If vision is used to achieve the goals of the organism, where a goal need not be a discrete state of the world, the system may not need to model the world at all before acting upon it.

Ecological Psychology and Robotics

Many of the papers in animate vision and robotics have made passing reference to the works of J.J. Gibson, but we would like to probe further into the relevance of his ideas (see Pickering (1992) for other points). Ecological psychology, as developed by Gibson (e.g., 1955, 1966, 1979), views animals and their environments as "inseparable pairs" that should be described at a scale relevant to the animal's behavior. So, for example, animals perceive the layout of surfaces, not the coordinates of points in space. A main tenet of the ecological approach is that the optic array, the pattern of light reflected from these surfaces, provides adequate information for controlling behavior without further inferential processing or model construction. This view is called *direct perception*: the animal has direct knowledge of, and a relationship to, its environment as a result of natural laws. How far into cognition perception plays a role is an open question, but minimally, the information involved in both perception *and* action could ground other, non-perceptual tasks. The strategy is to push natural law as far as possible into cognition, thus placing more constraints on the cognitive system.

The Gibsonian approach can be summarized in the idea that it is more desirable to put the animal in its environment than to put the environment in the animal. Rather than internally representing detailed knowledge of the world, animals detect and use information about it as it is required. This is the "fundamental hypothesis" of the ecological approach to vision:

> Optical structure specifies its environmental source and ... therefore, *mobile* organisms with *active* visual systems that can pick up this information will see their environments and suitably adjust their activity, if and when they detect that information, and only then (Turvey et al., 1981, p. 243, emphases ours).

Now, if we replace "mobile organisms" with "mobile robots," or more generally, "agents," this hypothesis is just as applicable to behavior-based robots as it is to animals. That is, sufficient information is available in the robot-environment interaction to control the robot's behavior without further inference or reconstruction. In addition, appropriate perception-action dynamics in the robot provide a non-inferential source of information upon which other aspects of computation (planning, mapping, reasoning, etc.) can be based and by which they can be limited.

Similar hypotheses might be made in regards to the other senses or sensors, but it is primarily vision that seems the most promising for unifying the fields of robotics and ecological psychology. Both would gain from such a union. The latter can provide insights into what kinds of information can control the actions of agents, that is, what ecological laws are at work for a given task; and the former can provide an experimental and demonstrative setting in which to test the viability of proposed control strategies and facilitate the discovery of new ones. The new robotics and the ecological approach complement each other well and both ultimately have the same concerns; thus, mobile robotics provides a promising test bed for ecological principles.

Optical Flow and Control Laws

A relevant case is the study of optical flow. As an observation point moves through the environment, the pattern of light reflected to that point changes continuously, creating *optical flow* (Lee, 1980; Gibson, 1958). Optical flow contains information about both the layout of surfaces and the motion of the point of observation. For example, if an observer is translating, the *focus of expansion* (FOE), or center of the radial flow pattern, specifies the observer's heading. If the observer is moving at a constant velocity, then the time-to-contact with a surface is given by the relative rate of expansion $\tau = r/v$, where τ is the "optic variable" *tau*-global (Tresilian, 1991; Lee, 1976), r is the visual angle between a point on the surface and the FOE, and v is the rate of change in this angle.

The observer's heading and time-to-contact are just two examples of information available in optical flow. One way an agent can use this information is by acting to achieve a certain *type* of flow. For example, to maintain ambient orientation, the type of optical flow required is no flow at all. If some flow is detected, then the agent should change the forces produced by its effectors (e.g., wings or legs) so as to minimize this flow, according to a *Law of Control*:

$$\Delta F_{\text{internal}} = f(\Delta \text{flow}) \qquad (1).$$

That is, the change in the agent's internal forces (as opposed to external forces like wind) are a function of the changes in the optical flow (here, from no flow to some flow).

Gibson (1958) described various control laws an animal might use for locomotion:

> ...to begin locomotion, therefore, is to contract the muscles so as to make the forward optic array flow outward. To stop locomotion is to make the flow cease.... To aim locomotion at an object is to keep the center of flow of the optic array as close as possible to the form which the object projects (1958, p. 187).

These types of rules have been noted by scientists studying the control of balance, steering, and braking in humans (Lee, 1976; Lee & Lishman, 1977; Yilmaz & Warren, in press; Warren, *et al.*, in press) and the control of flight in flies (e.g., Collet & Land, 1975; Reichardt & Poggio, 1976; Wagner, 1986a-b). Ambient orientation, or hovering, is controlled by minimizing the global optical flow: purely vertical flow (say, upward) will induce increased lift by the fly to minimize that flow (Srinivasan, 1977; Götz, *et al.*, 1979). Similarly, a fly in a rotating drum will produce a differential thrust with the two wings, tracking the drum by rotating about its own vertical axis (Collett, 1980a-b).

Warren (1988) proposed a set of control laws a fly might use for each of its major activities. For example, the laws of control for hovering in the face of vertical and horizontal flow, respectively, could be

$$\Delta U = (k/c)\Delta \dot{y} \qquad (2)$$

$$\Delta(F_L - F_R) = (k/c)\Delta \dot{w} \qquad (3),$$

where U is the amount of upthrust given by the two wings, (k/c) is the ratio of the drag constant to an optical scaling coefficient, $\dot{y}$ is the vertical component of the optical flow, F is the forward thrust given by a wing, and $\dot{w}$ is the horizontal component of the optical flow.

Which control laws govern the fly's behavior at any one time depend upon the goal, or "global action mode" of the fly (Warren, 1988): cruising, landing, foraging, pursuing conspecifics, etc. For each of these global action modes, objects in the environment will "afford" certain actions. Strictly speaking, the "affordances" of surfaces in the environment are constant for a particular animal (Gibson, 1979), but the global action mode determines which ones the fly uses. For example, while foraging, a fly will use a flower's affordance of nourishment and support, while avoiding all other surfaces. However, when tired, the fly might avoid flowers, but use the affordance of a resting place which large stationary objects will have. Once an action mode is adopted, the laws of control direct the actual output of the fly.

Ecological Robotics

We discuss below some work demonstrating that the control laws outlined in the previous section can be used successfully to control a mobile robot. We call this practice *ecological robotics*. It should be noted though that any mobile agent (be it biological or artificial) with a device to register the optical flow can use control laws like these to achieve its goals, with only the adjustment of some constants. Thus the study of optical flow for the control of action provides a domain in which experiments in two separate fields, ecological psychology and mobile robotics, can have direct relevance to each other.

Obstacle Avoidance

Our first work (Duchon & Warren, 1994) looked into control laws for the most crucial ability of a mobile agent: avoiding obstacles. The robot had a 12-inch base and a single camera with a 60° field of view placed about 75 cm off the ground. A fast patch-matching optical flow algorithm (Camus, 1994) provided a dense, robust flow field at 4 frames per second, allowing us to control the robot moving at a speed of 4 cm/s. Because we tested the robot in a tightly constrained office environment with poor lighting, we equipped the robot with a couple of "emergency" reflexes which would stop and turn the robot 90° when it got too dark (i.e., when no flow could be seen) and when a crash was immanent (τ<1 sec).

We investigated the performance of two control laws in this environment. In the Balance Strategy, the agent moves so as to equate the average magnitude of optical flow seen on each side of the FOE (as some bees do [Srinivasan, 1992]). In the Avoid-Closest Strategy the agent turns away from the point in the visual field with lowest τ.

We generally found the Balance Strategy to be more stable since it takes the entire field of view into account when determining the amount of rotation, whereas the Avoid-Closest Strategy is based only on a local region. For example, to go through an aperture, the Balance Strategy allowed the robot to head straight down the middle (where the flow would be equal on the two sides), whereas the Avoid-Closest Strategy required the robot to sequentially avoid one side or the other of the aperture until the sides were no longer in the field of view. However, both of the control strategies allowed the robot to wander and avoid obstacles successfully as illustrated in Figure 1. Dark regions under the tables and textureless chair backs caused problems for the robot, but its ability to avoid hands placed suddenly in its path demonstrates the utility of these control laws.

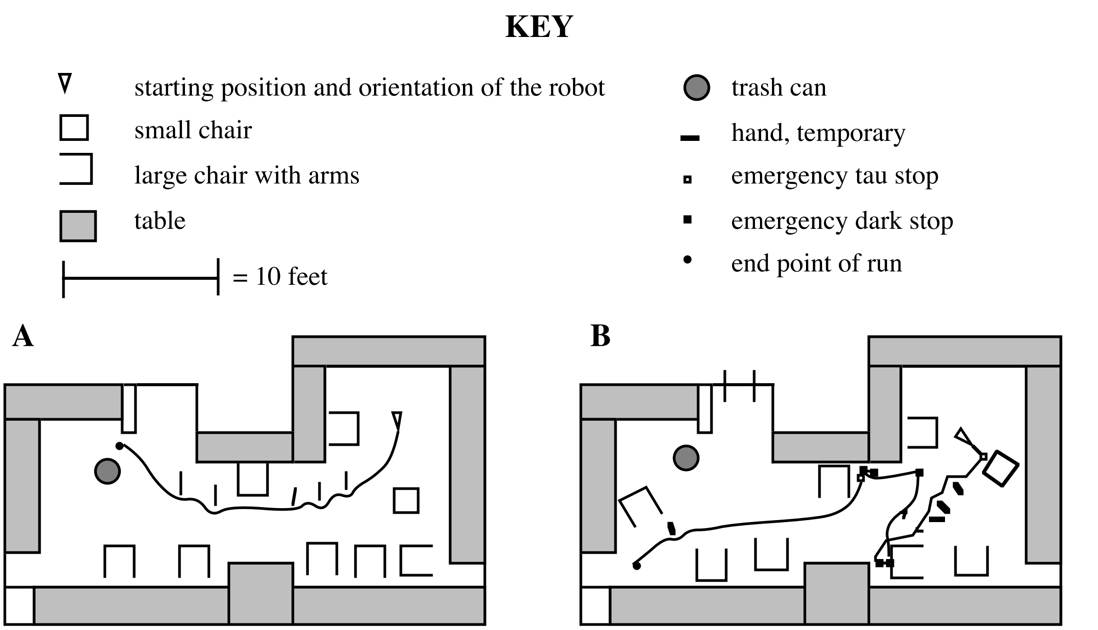

Figure 1. *Obstacle Avoidance*. **A:** *Balance Strategy*. The robot successfully avoids the hands placed in its path. The trial is stopped due to floor debris after 500 frames. **B:** *Avoid-Closest Strategy*. The robot is positioned just in front of a chair and the *tau* reflex makes it stop and turn 90° to the right. It avoids the hands and just misses a chair arm before running into a dark area. It has a problem again with the darkness under the tables, but the reflexes eventually point the robot in a favorable direction and it heads down the middle of the room towards the farthest corner. The trial is stopped due to cable lengths after 1000 frames.

At least two other groups have independently implemented similar ideas. Sandini *et al.* (1993) built a robot which would balance the flow seen from two cameras facing laterally. Coombs *et al.* (in press) have recently designed a robot with two cameras facing forward, one with a wide angle lens (115°) and one more foveal (40°), both of which are controlled with active gaze stabilization. Whereas they balance the *maximal* normals of the optical flow in the left and right *peripheral* fields, we take an *average* of the *entire* left and right fields relative to the FOE. Thus, the Balance Strategy used here can avoid a head-on collision with a well-textured wall from only 40 cm away (no such surfaces were present in the trials of Figure 1), because noise in the system breaks the symmetry and once a small difference in the flows is acted upon the difference becomes greater until the robot has completely avoided the wall. Also, no gaze stabilization is required since during a fast rotation, the amount of flow should be equal on the two sides, specifying no rotation, so optical flow due to the agent's own rotation is cancelled.

Other robotic implementations based on the insect literature (e.g., Cliff (1992), Franceschini *et al.*, (1992), and Sobey (1994)) have all used metric distance calculations and a mapping function for planning a (local) path or for tracking (cf. below). This information is totally absent here. Aloimonos (1992) and Nelson & Aloimonos (1989) took a similar approach to ours but required an additional intermediate representation of a "hazard map" based on normal flow to find the heading of the safest path.

The Game of Tag

Wandering around and not hitting things may help an agent avoid getting hurt, but survival requires more goal-directed behavior such as that exhibited in predator-prey interactions. A lion attacking a herd of zebras will chase the closest or slowest animal in the herd, and then use its claws to bring down the prey. The prey, for its part, must recognize that it is being attacked and make sure that the predator comes no closer. In humans, the children's game of tag is a (usually) gentler form of this interplay.

Our implementation of tag is slightly different from the typical game. We consider tag to be a global action mode, like foraging, and the aspects of tag like chasing and escaping to be subordinate action modes. Instead of the concept of *It*, there is only an *agent* and a *target*. While in the *watching* mode, the agent does not move until a moving *target* appears in its field of view. It then fixates the target by centering it in the field of view and tracks it throughout the trial. In our simplest implementation, no segmentation of the target is done. Instead, the control law for *fixating* is simply to turn to the right (increase force on the left) if more flow is seen on the right:

$$(F_L - F_R) = k\left(|\dot{w}_R| - |\dot{w}_L|\right) \quad (4).$$

While watching, this means that the robot will turn towards the side where the target appears and will continue to do so until the amount of flow is equal on the two sides of the field of view, i.e., it has fixated the target. This control law will be functional as long the target is the sole or fastest moving object (e.g., a rabbit in a field), the target is the closest object, all the moving objects are potential targets (e.g., a school of fish), or the motion signals are subsequently filtered (see Prokopowicz et al. (1994) for further discussion of how and when to use various kinds of information for tracking—though we would argue that motion signals can still be used when the agent is moving).

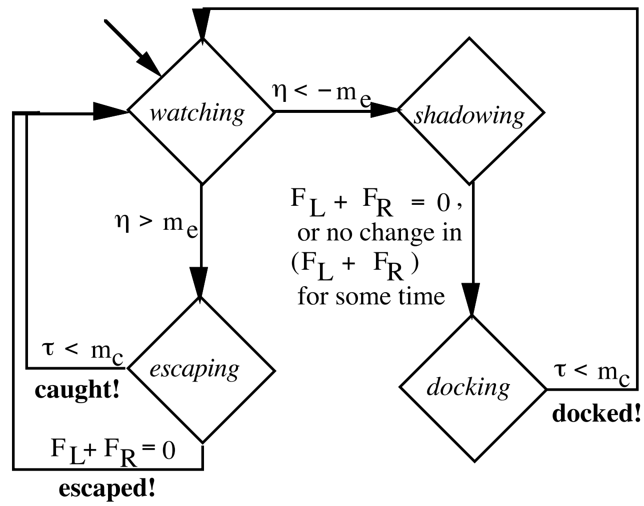

Figure 2. *The action modes for tag and the transitions between them. Fixating* takes place continuously throughout the game. A **miss!** was recorded if the target fell outside the robot's field of view. The transition conditions and control laws for each mode are described in the text.

If the target withdraws from the agent (i.e., the optical flow has an focus of contraction, detectable by $1/\tau = \eta < -m_e$, where m_e is a "margin value") then the agent chases it. If the target approaches the agent (i.e., $\eta > m_e$) then the agent enters the *escape* mode, backing up until it is caught or successfully escapes. Chasing is defined as having two parts: a) the *shadowing* mode in which the agent matches the speed of the target; and b) the docking mode in which the agent makes a controlled approach (rather than a hard attack) with the target. Shadowing is achieved by

$$\Delta(F_L + F_R) = -k\eta \quad (5)$$

i.e., increasing force if the target is escaping and decreasing force if the agent is gaining on it. Escaping is essentially the same except that the agent needs to continue to increase the distance from the target at all times, which can be done by adding the term $-k\mu$ to the right-hand side of equation (5), where μ is the minimum amount of acceleration from the target. The agent exits the escaping or shadowing modes if it has stopped accelerating.

The second aspect of chasing, *docking*, is achieved by

$$\Delta(F_L + F_R) = k\left(\dot{\tau} + 0.5\right) \quad (6)$$

where $\dot{\tau}$ is the derivative of τ, and can be used to control deceleration prior to contact. If $\dot{\tau}$ is kept equal to -0.5, the observer will just touch the target (Lee, 1976). Docking is complete when τ is below a certain margin value, m_c, which can also be used to register that it has been caught if it is escaping. Once the escape or dock is complete, the agent stops and the process begins again—the agent waits for a target to come into its field of view. Figure 2 gives a summary of these control laws. (Further discussion can be found in Warren (1988)).

In the office environment of Figure 1, we videotaped a series of escapes and chases. The experimenter brought a 4×12 inch textured target into the robot's field of view and moved it away from (or towards) the robot at an appropriate rate and with a few changes in direction. Even without explicit segmentation of the target, the robot successfully chased down the target 70% of the time, with some chases lasting 3-4 meters. Escaping was more difficult since the robot had to move backwards while facing the target and only 40% of the escapes stopped soon after the target stopped approaching.

We also implemented these controls laws with segmentation based on a simple difference-of-boxes operator acting on the horizontal dimension of the optical flow, but found that the percent of successful chases did not increase. This may have been due to the crude nature of the segmentation algorithm. Animals undoubtedly have more sophisticated means of segmentation which would include other properties of the target (e.g., color, shape, size, and type of internal motion). Nevertheless, segmentation is hard and as was seen with the obstacle avoidance algorithm, reliance on precise segmentation can lead to less robust performance. Psychophysical experiments in our lab (Warren & Saunders, in press) indicate that humans do not segment the scene before determining their heading either. Further work is underway to investigate the extent to which 3-D models of the environment are used by humans when navigating under circumstances similar to the robots here.

Conclusion

We have discussed behaviors like obstacle avoidance and the game of tag which can be produced in a robot with no reconstruction of the visual scene (Aloimonos, 1993). At a minimum, this work points to an approach eschewing a central model in favor of a tighter binding between action and perception. This methodology has been explored by a number of robotics researchers (e.g., Aloimonos, 1992; Brooks, 1991b; Coombs, *et al*., 1995; Horswill, 1993; Pfeifer and Verschure, 1993; Sandini *et al*., 1993) and has even produced higher-order behaviors like planning (e.g., Mataric, 1992; Meeden, 1994). The similarities between these approaches and ours based independently on Gibsonian ideas suggests that the application of the theories and results from forty-five years of ecological psychology will surely enhance this endeavor.

Our work also ties in with recent physiological studies and lesion cases (Milner & Goodale, 1993) that suggest separate "what" and "how" pathways in the brain. The lesion cases have indicated a difference between knowing what an object is, and knowing how to maneuver it. This change of emphasis is also reflected in some philosophical approaches to knowledge whereby to "know that" first requires one to "know how" (Bechtel, 1990; Ryle 1949). Our robot does not need to know what an object is in order to avoid it, nor identify a target before knowing how to control its escape. In essence, we have implemented a simple "how" pathway. But, since an approaching conspecific may afford mating as well as escape, it is important that all the affordances of an object be recognized and one of them acted upon. Neural networks would be an ideal means of satisfying the many soft constraints (affordances) of an object and choosing a single output (action mode) (see also Brooks' subsumption architecture [1991b] and Pfeifer & Verschure's distributed adaptive control [1993]). In any case, knowledge of the affordances of an environment provides a basis for a choice of action, and that action, once chosen, can be controlled without a central model of the world. Such procedural, functional knowledge seems necessarily prior to more abstract, declarative knowledge.

Finally, the fact that these control laws are essentially universal for mobile agents with perceptual systems capable of detecting optical flow means that they can be investigated in insects, animals, humans and robots. We are beginning an interactive approach with the last two kinds of agents. Robotic modeling helps us determine the plausibility of control laws that have been hypothesized for biological agents, and from psychophysical studies we hope to find new control laws useful in a robot. The study of control laws based on optical flow thus provides a unique opportunity for cognitive scientists, computer scientists and engineers to work together, solving the same problems.

Acknowledgments

Tom Dean kindly provided the robots and cameras. Ted Camus' work on real-time optical flow made this entire project possible. Jak Kirman, Moises Lejter, and Jon Monsarrat were very helpful in getting the robots and computers working. This work was supported by a National Science Foundation Graduate Research Fellowship to the first author.

References

Aloimonos, Y. (1992). Is visual reconstruction necessary? Obstacle avoidance without passive ranging. *J. of Robotic Systems*, 9(6),843-858.

Aloimonos, Y. (1993). Introduction: Active Vision Revisited. In Y. Aloimonos (ed.), *Active Perception*, (pp. 1-18). Hillsdale, NJ: Lawrence Erlbaum Associates.

Aloimonos, Y. & Rosenfeld, A. (1991). Computer Vision, *Science*, 253, 1249-1254.

Ballard, D. (1991). Animate vision. *Artificial Intelligence*, 48, 57-86.

Ballard, D.H. & Brown. C.M. (1992). Principles of animate vision. *CVGIP: Image Understanding*, 56(1), 3-21.

Bechtel, W. (1990). Multiple level of inquiry in cognitive science. *Psychological Research*, 52, 271-281.

Brooks, R.A. (1991a) New approaches to robotics. *Science*, 253, 1227-1232.

Brooks, R.A. (1991b). Intelligence without representation. *Artificial Intelligence*, 47, 139-160.

Camus, T. (1994b). Real-time optical flow. Doctoral dissertation and technical report CS-94-36. Providence:

Brown University.

Cliff, D. (1992). Neural networks for visual tracking in an artificial fly. In *Proc. of the 1st European Conf. on Artificial Life* (pp. 78-87). Cambridge, MA: MIT Press.

Collett, T.S. (1980a). Some operating rules for the optomotor system of a hoverfly during voluntary flight. *J. of Comparative Physiology*, 138, 271-282.

Collett, T.S. (1980b). Angular tracking and the optomotor response: An analysis of visual reflex interaction in a hoverfly. *J. of Comparative Physiology*, 140, 145-158.

Collet, T.S. & Land, M.F. (1975). Visual control of flight behavior in the hoverfly, *Syritta pipiens* L. *J. of Comparative Physiology*, 99, 1-66.

Coombs, D., Herman, M., Hong, T., & Nashman, M. (in press). Real-time obstacle avoidance using central flow divergence and peripheral flow. In *Proc. of the 5th Int. Conf. on Computer Vision*.

Duchon, A.P. & Warren, W.H. (1994). Robot navigation from a Gibsonian Viewpoint. In *IEEE Int'l Conf. on Systems, Man, and Cybernetics* (pp. 2272-2277). Piscataway, NJ: IEEE.

Franceschini, N., Pichon, J.M., and Blanes, C. (1992). From insect vision to robot vision. *Phil. Trans. R. Soc.of Lond. B*, 337, 283-294.

Gibson, J.J. (1950) *The Perception of the Visual World*. Boston: Houghton Mifflin.

Gibson, J.J. (1958). Visually controlled locomotion and visual orientation in animals. *British J. of Psychology*, 49(3), 182-194.

Gibson, J.J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.

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

Götz, K.G., Hengstenberg, B., & Biesinger, R. (1979). Optomotor control of wing beat and body posture in Drosophilia. *Biological Cybernetics*, 35, 101-112.

Horswill, I.D. (1993). Specialization of perceptual processes. Doctoral Dissertation. Cambridge: MIT.

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. (1980). The optic flow field: The foundation of vision. *Phil. Trans. R. Soc. of Lond. B*, 290, 169-179.

Lee, D.N. & Lishman, R. (1977). Visual control of locomotion. *Scandinavian J. of Psychology*, 18, 224-230.

Mataric, M. (1992). Integration of representation into goal-driven behavior-based robots. *IEEE Trans. on Robotics and Automation*, 8(3), 304-312.

Meeden, L.A. (1994). Towards planning: Incremental investigations into adaptive robot control. Doctoral Dissertation. Bloomington: Indiana University.

Milner, D.A. & Goodale, M.A. (1993). Visual pathways to perception and action. In Hicks, T.P., Molotchnikoff, S. & Ono, T. (eds.) *Progress in Brain Research*. Volume 95. (pp. 317-337). North-Holland:Elsevier Science Publishers.

Moravec, H.P. (1981). Rover Visual Obstacle Avoidance. In *Proc. of the 7th International Joint Conference on Artificial Intelligence* (pp. 785-790).

Nilsson, N.J. (1984). Shakey the robot. *SRI International Technical Note*, No. 325.

Pfeifer, R. & Verschure, P. (1993). Designing efficiently navigating non-goal-directed robots. In *Proc. of the 2nd Int'l Conf. on the Simulation of Adaptive Behavior* (pp. 31-39). Cambridge, MA: MIT Press.

Pickering, J. (1992). The new artificial intelligence and biological plausibility. In: *Studies in Perception and Action II: Posters Presented at the VIIIth Int'l Conf. on Event Perception and Action* (pp. 126-129). Hillsdale, NJ: Lawrence Erlbaum.

Prokopowicz, P.N, Swain, M.J., & Kahn, R.E. (1994). Task and environment-sensitive tracking. In *Proceedings of the IARP/IEEE Workshop on Visual Behaviors* (pp. 73-78). Piscataway, NJ: IEEE.

Reichardt, W. & Poggio, R. (1976). Visual control of orientation behavior in the fly. *Quarterly Reviews of Biophysics*, 9 (3), 311-375.

Ryle, G. (1949). *The Concept of Mind*. New York: Barnes and Noble.

Sandini, G., Santos-Victor, J. Curotto, F. & Garibaldi, S. (1993). Robotic Bees. In *Proc. of IROS-93*, Yokohama, Japan.

Sobey, S. (1994). Active navigation with a monocular robot. *Biological Cybernetics*, 71, 433-440.

Srinivasan, M.V. (1977). A visually-evoked roll-response in the housefly. *J. of Comparative Physiology*, 119,1-14.

Srinivasan, M.V. (1992). How bees exploit optic flow: Behavioural experiments and neural models. *Phil. Trans. R. Soc. Lond. B*, 337, 253-259.

Tresilian, J.R. (1991). Empirical and theoretical issues in the perception of time to contact. *J. of Experimental Psychology: HPP*, 17, 865-876.

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, 237-304.

Vera, A.H. & Simon, H.A. (1993). Situated Action: A symbolic interpretation. *Cognitive Science*, 17, 7-48.

Wagner, H. (1986a). Flight performance and visual control of flight of the free-flying housefly (*Musca Domestica* L.) II. Pursuit of targets. *Phil. Trans. R. Soc. of Lon. B*, 312, 553-579.

Wagner, H. (1986b). Flight performance and visual control of flight of the free-flying housefly (*Musca Domestica* L.) III. Interactions between angular movement induced by wide- and small-field stimuli. *Phil. Trans. R. Soc. of Lon. B,* 312, 581-595.

Warren, W.H. (1988). Action modes and laws of control for the visual guidance of action. In Meijer, O.G. & Roth, H, (eds.) *Complex Movement Behaviour: The motor-action controversy* (pp. 339-380). North-Holland: Elsevier Science Publishers.

Warren, W.H., Kay, B. & Yilmaz, E. (in press) Visual control of posture during walking: Functional specificity. *J. of Experimental Psychology: HPP*.

Warren, W.H. & Saunders, J. (in press) Perception of heading in the presence of moving objects. *Perception*.

Yilmaz, E. & Warren, W.H. (in press) Visual control of braking: A test of the tau-dot hypothesis. *J. of Experimental Psychology: HPP*.