Phototaxic Foraging of the Archaepaddler, a Hypothetical Deep-Sea Species

R. J. V. Bertin
Laboratoire de Physiologie de la Perception et de l'Action
Collège de France / C.N.R.S.
11, place Marcelin Berthelot
75005 Paris, France
bertin@cdf-1ppa.in2p3.fr

W. A. van de Grind
Neuroethology Group
Department of Comparative Physiology
Utrecht University
and
Helmholtz Instituut
School for Autonomous Systems Research
Padualaan 8
3584 CH Utrecht
The Netherlands
W.A.vandeGrind@bio.uu.nl

Abstract An autonomous agent (animat, hypothetical animal), called the (archae) paddler, is simulated in sufficient detail to regard its simulated aquatic locomotion (paddling) as physically possible. The paddler is supposed to be a model of an animal that might exist, although it is perfectly possible to view it as a model of a robot that might be built. The agent is assumed to navigate in a simulated deep-sea environment, where it forages for autoluminescent prey. It uses a biologically inspired phototaxic foraging strategy, while paddling in a layer just above the bottom. The advantage of this living space is that the navigation problem—and hence our model—is essentially two-dimensional. Moreover, the deep-sea environment is physically simple (and hence easy to simulate): no significant currents, constant temperature, completely dark. A foraging performance metric is developed that circumvents the necessity to solve the traveling salesman problem. A parametric simulation study then quantifies the influence of habitat factors, such as the density of prey, and body geometry (e.g., placement, direction and directional selectivity of the eyes) on foraging success. Adequate performance proves to require a specific body geometry adapted to the habitat characteristics. In general, performance degrades gracefully for modest changes of the geometric and habitat parameters, indicating that we work in a stable region of "design space." The parameters have to strike a compromise between, on the one hand, the ability to "fixate" an attractive target, and on the other hand, to "see" as many targets at the same time as possible. One important conclusion is that simple reflex-based navigation can be surprisingly efficient. Additionally, performance in a global task (foraging) depends strongly on local parameters such as visual direction tuning, position of the eyes and paddles, and so forth. Behavior and habitat "mold" the body, and the body geometry strongly influences performance. The resulting platform enables further testing of foraging strategies or vision and locomotion theories stemming either from biology or from robotics.

Keywords
hypothetical animal, two-light experiment, phototaxic foraging, visuomotor behavior, foraging performance, animats

1 Introduction

Upon shading or switching off the light, the "dog" can be stopped immediately, but it will resume its course behind the moving light so long as the light reaches the condensing lenses in *sufficient intensity*. Indeed, it is more faithful in this

> respect than the proverbial ass behind the bucket of oats. To the uninitiated the performance of the pseudo dog is very uncanny indeed.

This is a quote from an article in the *Electrical Experimenter* of September 1915 (quoted in [22, pp. 68–69]), describing the "Orientation Mechanism" built by the "well-known inventor, Mr. John Hays Hammond, Jr." It is probably the first example of what is now called artificial life, far predating Grey Walter's description of similar, though monocular, machines in the 1950s [32, 33]. Hammond appears to have been quite aware of the possibilities of making practical use of such a navigation mechanism, because he applied it to the "Hammond Dirigible Torpedo."

Loeb [22] was one of the firsts to analyze visually guided animal navigation, the topic of this article, in bold mechanistic terms. He rejected the century-old idea that phototropic navigation is based on the animal's feelings of "love" or "hate" for light. According to Loeb the symmetry plane of the animal's body is the basic orientation reference. A frontal sensor on one side of this plane could innervate a caudal motor appliance on the opposite side of the symmetry plane, leading to an automatic orientation of the symmetry plane so as to contain the light source, presuming the presence of a single source. These considerations lead quite naturally to the question of how an animal would orient in the presence of two lights. If the two lights are placed at equal distances from the symmetry plane on a line perpendicular to that plane one has the equivalent of the donkey's dilemma between two sheaves of hay. Kühn [20] and Fraenkel and Gunn [12] modified and expanded on Loeb's ideas and radically changed the terminology (e.g., Loeb's tropism is now called a taxis). Countless animal species were subjected to the two-light experiment [12, 20, 22, 25] and it therefore seems like a natural point of departure for a study on visually guided navigation, such as ours. Therefore we developed a bilaterally symmetric class of hypothetical animals, the paddlers, with two eyes in the head region and two paddles in the tail-region, placed symmetrically relative to the midsagittal plane, which is Loeb's symmetry plane. To keep things simple and yet "natural," we assume that these animals live in the deep sea (no turbulence, no background light, constant temperature, etc.) and feed on another hypothetical species, glowballs. Of course glowballs are autoluminescent, so they look like submerged light bulbs. More details are given later on.

The two-light experiment is of interest because it has been carried out on many animals and because it gives perhaps the simplest and most basic information on orientation and navigation in an environment with more than one target. However, one might wonder why animals would bother to move toward or away from light sources in the first place. Our paddlers have a clear and vital reason: to fill their stomachs and survive. In such a situation it would be reasonable to expect that evolution leads to more sophisticated navigational strategies and better methods of choosing the next target in a larger collection of food items. We were therefore interested in two basic questions. First, if a very simple paddler, almost built to the specifications of Loeb, behaves as expected in the two-light (two-glowball) situation, how does it behave in a living space with many glowballs? Second, how can or should the inevitable shortcomings be remedied? In other words, what are the simplest conceivable improvements on Loeb's scheme? To study these questions quantitatively one has to solve a number of auxiliary problems. First, how do we quantify foraging success? Second, what is the influence of body geometry on foraging success?

Phototaxis is the "visual version" of *tropotaxis*, a kind of autonomous closed-loop navigation in which sensory excitation in the nervous system is kept (left-right) symmetric by motor actions governed by the same sensory information ([20], following [22]). This principle is nicely illustrated in [6]. Tropotaxis of chemical nature (*chemotaxis*) is demonstrated in [2]. We will use the term *phototaxic navigation* for visually guided

tropotaxis based on a bisensor system (two eyes) with a directional sensitivity tuning of each of the two sensors (eyes). A Gaussian weighting function gives each eye a specific direction of maximal light sensitivity (the line of sight) flanked by regions in which sensitivity declines smoothly, as it does in many bisensor systems [25]. If the width parameter of this weighting function is made very large, one has the Loeb system, without visual direction tuning. Therefore we add this "innovation" of visual direction tuning right from the start and regard the archaepaddler L (L for Loeb) as a limiting case of the more general archaepaddler B (B for binocular).

There are several levels of detail in which acting animals can be modeled. An exhaustive review of biologically inspired models of navigation is given in [31]. Here we can only give a (necessarily incomplete) overview of related or relevant studies. In general, more complex behaviors (such as the learning of complicated visuomotor coordination, e.g., [9], or studies focusing on social interactions, e.g., [30]) are best modeled at a higher, more schematic level of description, with more detailed, biophysical aspects (such as those of locomotion) approximated or taken for granted. This achieves a trade-off, losing some realism in the parts of lesser interest but preserving it in the other parts, all the while keeping the computing expenses at acceptable levels. Typically, such models will generate a certain "do-this-or-that" command through the mechanisms studied, which is then executed. Similarly, studies that are mostly concerned with low-level control of locomotion, as in the model of the lamprey's swimming central pattern generator described in [11], may disregard the fine(r) aspects of sensory perception, directly providing the modeled motor system with precise commands that are then "executed" with simulated physical realism, or even in physical reality (as is the "stickbot" described in [10]). However, the animal's actions are interactions with its environment, and as such they directly influence its sensory input, which in turn is the interaction of the environment with the animal. Therefore studies in which sensorimotor behavior is modeled at a more detailed level of description, for instance, in the context of evolving models of sensorimotor systems that can at some point be "implemented" as real robots (e.g., [8]), do explicitly model biophysical aspects of sensory perception and motor behavior. The same applies to studies aiming at realistic animation and/or interaction of behaving animals, such as the artificial fishes ([29, 30]) that learn to perform a variety of motor skills, giving rise to "experimental results" of impressive photorealistic quality.

Here, we simulate the physical processes of locomotion in sufficient detail to view it as representative of locomotion in real animals. Our general interest motivating the research presented here lies in getting a feeling for the kind of evolutionary constraints, obstacles, or problems an animal species might encounter during its evolution, focusing on individual, elementary behaviors such as foraging and not so much on more complex behaviors such as social interactions. We purport to take inspiration from, but also return inspiration to, neurobiology and neuroethology. This requires a high degree of explicitness regarding the mechanisms of vision, locomotion, and navigation. This aim contrasts with, for example, the approach described by Terzopoulos [29], who uses computer vision algorithms, a "perceptual oracle" with direct access to a graphics rendering machine, a color indexing algorithm by Swain [28], and scalar actuation functions for muscle innervation. The result is exciting and would allow a wonderful kind of virtual ethology. Nevertheless, the approach is radically different from our aim to stick as closely as possible to known neuronal mechanisms and to principles likely to be implemented in even the lowliest of animals. Such animals are, for example, known to be able to recognize food objects without having a sophisticated color analysis system. We use only those simple operations and components that are universally used in neurobiological modeling (leaky integrators as membrane mimics, subtractive and shunting inhibition, thresholds and summation). The "processing principles" are of

a simple (local) type and we then obtain navigation behavior as emergent property. We are interested in what emerges at the behavioral level from plausible and simple visuomotor interactions. However, the behavioral level is quantified in terms of a fitness parameter, which allows us to reverse the causality and see how the local properties (of body geometry, eye structure, neuronal interaction) are molded by selection (here only based on hand picking species on the basis of foraging quality). We are interested in this two-way interaction between local properties and global foraging behavior with nothing in between but the simplest known sensorimotor networks.

It would of course be nice to simulate a complete evolution of the paddlers. In that case, one would merely need to specify the environment, a simple ancestral species, and a selection criterion and then wait and see what evolves. The tools to perform such studies are available ([8, 26, 27]) and have been shown to lead to virtual creatures with behavior not unlike that of our paddlers [26]. However, given the level of detail in our model, and the number of individuals and generations necessary to achieve evolutionary adaptation, this is not a very feasible option.[1] Therefore we resorted to tinkering in the sense of the evolutionary process [19]. With our criterion of foraging success, as specified later on, we carried out a parametric study of paddlers. This enabled us to fix all parameters so as to optimize foraging success. Further improvement then requires some innovation. The boundary conditions for choosing an innovation are a) that it is known to exist in some animal and b) that it seems like a small step relative to the existing design. For example, after developing the simple archaepaddler L and studying its behavior in a multi-target environment like ethologists study real animals, we conclude the animal needs some mechanism for course stabilization such as a mechanism to focus on one target even when there are many in view. Two well-known principles were implemented to achieve that goal, namely directional weighting (tuning) of the visual receptor layer and giving binocular information additional influence to emphasize straight-ahead relative to course changes. These innovations lead to a new species, the archaepaddler B, which proves to perform quite well in complex multi-target environments. It is in turn the starting point for the development of a more sophisticated navigator described elsewhere [5], a shallow-water paddler that can handle a wide range of background light intensities and both positive and negative contrast of the target. These adaptations are necessitated by migration to shallower water.

2 The Archaepaddler

Figure 1 depicts the archaepaddler (or simply *paddler*, for the purposes of this article), a hypothetical animal with a circular body of radius R. Apart from a pair of eyes, a mouth, a "catch region," paddles for locomotion, and a pair of pectoral fins (*pectfins*) to help in course stabilization, there is a simple nervous system (NS), described below. The first problem we now face is to choose a body structure for the paddlers that is compatible with Loeb's ideas but at the same time specifies all the parameters. Where do we place the eyes and paddles, and at what angle relative to the plane of symmetry? How wide do we make the visual field of the eyes and in what direction should the visual weighting be maximal (the direction of the line of sight)? Because all these factors are under survival pressure, we cannot fix them a priori. The only option left is to introduce a list of structural parameters and do a parametric study. The eyes are positioned at an angle α relative to the body center and the rostro-caudal body axis and have a visual field of φ degrees centred on their eye axis, which makes an angle of β degrees with the rostro-caudal axis. For the purposes of this article the exact

1 To give an idea, on a 150 MHz SGI R4400 Indigo^2 workstation, simulating three paddlers for 4,000 time units (i.e., collecting one datapoint on the graphs presented in the figures) takes between 10 and 20 minutes.

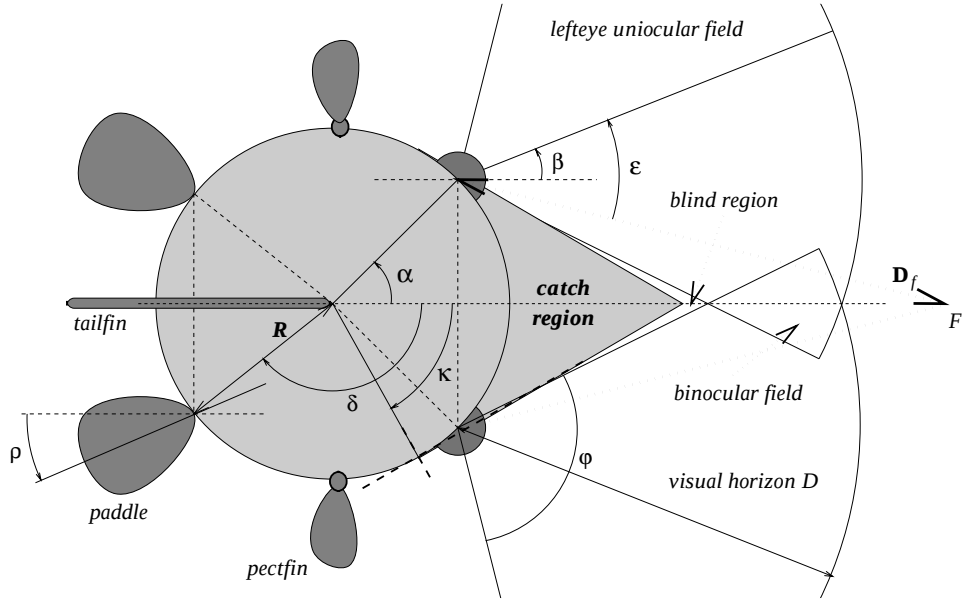

Figure 1. A paddler. The indicated parameters are explained in the text.

anatomy of the eyes (lens eye or compound eye) does not matter, provided it has some directional selectivity.

It is known that light of the dominant autoluminescent wavelengths (which humans would call *blue-green*) can be seen from a maximum distance of 10–16 m in the deep sea [21]. This "effective visual horizon" is the result of absorption and scattering by the water and the particles suspended therein. As an approximation, to be validated below, we therefore introduce an explicit visual-horizon parameter in the simulations, called D (maximum visual distance).

The retina consists of a one-dimensional array of $N = 80$ identical photoreceptors with identical, nonoverlapping acceptance angles. To give the eyes a best visual direction the contributions of these receptors to an eye's overall response are weighted as indicated in Figure 2. The retinal weighting function (RWF) has a Gaussian shape with halfwidth σ and an optimum direction called the visual axis. The visual axis makes an angle of ε degrees with the eye axis (Figure 1) in such a way that an overlap of the left and right fields of view arises and thus that binocular vision can be favored over monocular vision (see below). The visual axes cross in the "fixation point"[2] F on the paddler's longitudinal body axis (see Figure 1); the distance from the center of the eye to F is called the fixation distance D_f. A fixation distance of $D_f = \infty$ thus denotes a visual axis parallel to the paddler's longitudinal axis ($\varepsilon = \beta$).

The eye needs separate photoreceptors (or receptive fields) to resolve separate targets (i.e., glowballs), but in the archaepaddler, information about targets is pooled early in the visual system. Each eye's output is a visual-direction-weighted estimate of the amount of light reaching the eye from the corresponding part of the world. More precisely, it is the weighted sum of all photoreceptor outputs, divided by the number of supraliminally activated receptors. This principle of directional tuning of the bisensors

2 Just like we use the term "fovea" (the direction of maximal sensitivity) as an analogue to its standard meaning as "the retinal region of maximal resolution," "fixation" in this context does not imply that the paddler is actually capable of fixating a given point. Response to a target moving along the paddler's longitudinal axis is strongest at the fixation point.

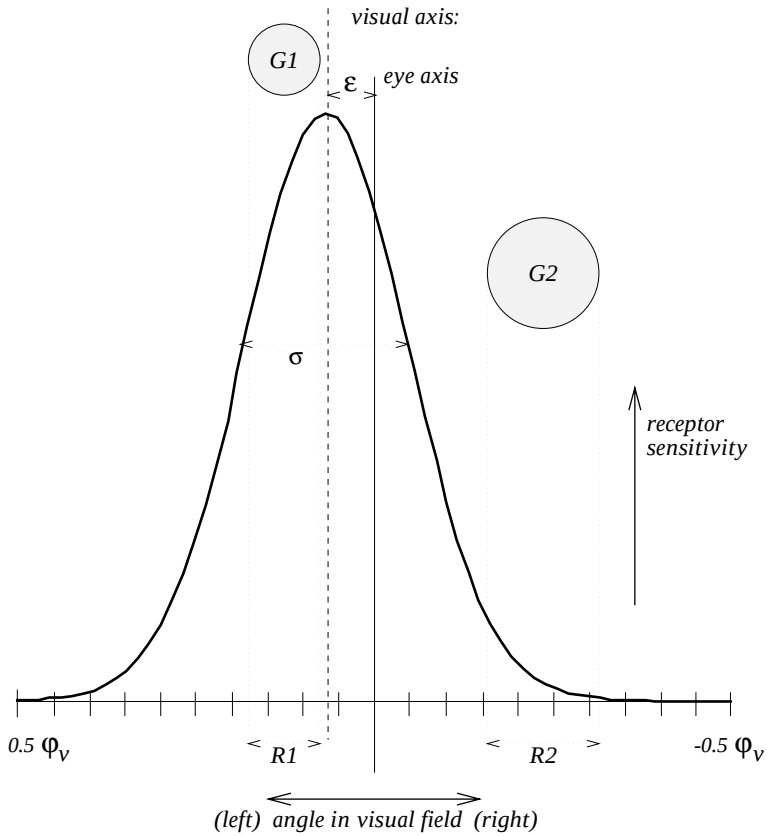

Figure 2. Retina and visual field of the paddler's eye in polar representation. The receptive fields that make up the retina can, as is indicated, consist of a number of photoreceptors projecting onto a set of ganglion cells that measure certain local properties of the visual field. In the present article, the receptive fields consist of only one photoreceptor. Due to the RWF, a Gaussian receptor-sensitivity profile, the response (R1) to a more distant (or smaller) glowball (G1) can be higher than the response (R2) to a closer or larger glowball (G2), if G1 is more "straight-ahead" than G2. This is similar to the well-known decrease of resolution with eccentricity.

can, depending on the choice of parameter values, facilitate course stabilization (or a crude form of fixation/selective attention—potentially useful capabilities) with a preference for straight-ahead or small yaw angles, independent of the exact mechanism of course determination. In fact, directional sensitivity improves performance (the reliability of finding a source) even in one-sensor organisms [18]. The influence of the weighting function can be seen intuitively from Figure 2 by considering a small glowball *G1* close to the optimal direction and a larger glowball *G2* in a more peripheral direction. The size and intensity of a glowball's retinal image depends on the size and distance of this target. The eye's response is proportional to the size and intensity of the retinal image and depends on retinal position as specified by the retinal weighting function (RWF). Thanks to the weighting, *G1* gives a stronger response. Thus if *G1* is seen by one eye and *G2* by the other eye the smaller course change necessary to catch *G1* will win out over the larger course change to catch *G2*, even though *G2* is bigger (or closer). The parameters of this weighting function allow us to play with this type of preference. Obviously if the halfwidth is large, directional tuning disappears, and *G2* in the example might win and the animal might more easily change course to catch bigger or closer prey at higher eccentricities. The choice of this parameter is one of the

differences between the L and B versions of the archaepaddler. The other difference is that the B paddler has an additional parameter ϖ ($\varpi = 0$ in the archaepaddler L), the binocular facilitation parameter that also plays a role in course stabilization (see later).

Summarizing, the eye position, eye axis, and field of view (parameters α, β, and φ) determine which portion of the outside world is sampled. The visual axis (ε) and RWF halfwidth (σ) determine how it is sampled, that is, how the sensitivity is distributed over the sampling area.

2.1 The Nervous System

The nervous system (NS; Figure 3) converts the eye-responses into the appropriate swim commands. It is a feed-forward network with analog components (neurons). Communication between the neurons is represented by real numbers (range [0,100]) that stand for spike frequencies. There are three types of synapses: excitatory (addition), inhibitory (subtraction), and shunting inhibition. Shunting inhibition is implemented by dividing input x by $(i + 1)$, where i is the inhibitory signal. The neurons' output is a linear function of the net input, clipped between the minimal and maximal "firing frequencies" mentioned above, except for the two leaky integrators introduced further down.

The NS is inspired partly on specific, real biological systems (e.g., the paddle controllers, see below), and partly on general principles found throughout the animal kingdom. It has three distinct centers. The *normalizing center* remaps the eye responses onto a constant interval by dividing each eye response by the sum of both eye responses. Responses from the binocular regions of the eyes are pooled and receive additional (ϖ) weight in the B paddler, but not in the L version. The *visuomotor center* maps the visual signals onto motor command and control. Two leaky integrators (modeled as first-order low-pass filters with time constant τ, generally taken to be reasonable models of neuronal transfer functions) filter out fast fluctuations and serve as a short-term memory for direction of movement, another course stabilization mechanism. A mutual inhibition between the two leaky integrators, which is suppressed in the presence of visual information, amplifies the last maneuver that started before visual information ceased, by silencing the weaker of the two leaky integrators. When the paddler leaves a "patch" of glowballs, entering an empty region, this amplification of the last movement (i.e., a turn to the left, a turn to the right, or—theoretically—a movement straight-ahead) serves to redirect the paddler to the patch it just left. In the absence of visual information, the paddler's movements are determined by a *wander controller*. This is a network of mutually inhibiting neurons with stochastic spontaneous activity. It generates swimming bouts of random speed and direction. During visually guided behavior it is inhibited by a neuron integrating the responses of both eyes. The same neuron gates the mutual inhibition between the two leaky integrators. A *motor center* uses the signals from the visuomotor center to create swim commands.

2.2 Locomotion and Its Control

The paddles are stiff appendages positioned as indicated in Figure 1. Movement of the paddles is controlled by two *paddle controller* networks. For the experiments described in this article, simplified *linear* paddle controllers were used. These controllers linearly map the central swim command onto a paddling force (thrust).

A speed control system regulates the paddler's actual swimming speed. A full description of this system is outside the scope of the current article: Here a general description of its functional principles will suffice. The speed control system derives the required speed from the central swim commands and compares it with the actual speed, setting the gain to the paddle controllers accordingly. If the actual speed is too large, the pectfins are tilted accordingly to brake. Actual speed is measured from deflection

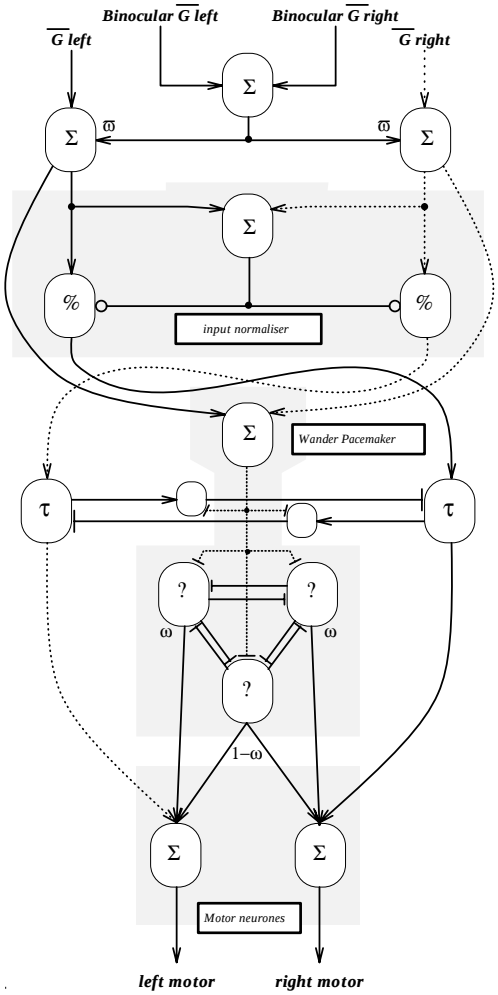

Figure 3. The visuomotor network of the paddler's nervous system. The symbols Σ, X, $\gg$ and ? indicate respectively a summating, a correlating, a threshold, and a "random" neuron; τ a leaky integrator neuron, and % a normalizing (dividing) neuron. The symbol $\rightarrow$ indicates an excitatory synapse, $\dashv$ an inhibitory synapse and $-\circ$ a shunting inhibitory synapse. The network is explained in the text.

(amount and direction) of a contralateral pair of hair receptors caused by the water flow around the body. The direction of movement is labeled-line encoded: Rostro-caudal and caudo-rostral water flow activates different channels. Required speed is similarly represented in swim direction/speed vectors. The gains to the paddle controllers, and possibly the brake commands, are determined using contralateral combinations of the required and actual speed information. The speed control system compensates for inertia, and also for the slower acceleration at low thrust (see below). Especially the paddler's turnability benefits from this. It also enables the paddler to swim forward by using its tail fin (much like a gondolier propels his gondola) when the paddles can only generate rotational thrust (as happens when the generated thrust is parallel to the body).

A more detailed description of the locomotor system—including a model of oscillatory paddle controller networks based on [34, 35]—and a motivation of the design are given in [4] and will be the subject of a future publication. The linear mapping used

in the current article is a good approximation of the mapping performed by the more complex oscillatory paddle controller.

A summary of the physics of the paddler's locomotion is given in Appendix A. Although some aspects of this design of the locomotor system might seem arbitrary at first glance, there really is not much freedom in designing such a system given the physical constraints and the goals of cruise control.

We will now turn to a description of our index of performance, the environment, and our simulation techniques and introduce the (default) values used for the different parameters of the paddler, its environment, and the experiments. In the experiments presented below we concentrate on the most influential parameters of the paddler and its habitat.

3 Methods

3.1 Performance Index

The following experiments address the influence of several of these parameters on the paddler's overall performance, as quantified in terms of a performance index $\wp$. The value of this index mirrors the "fitness" of the corresponding individuals. An exact performance index would relate the paddler's route to the optimal foraging route, which is of course found by solving the Traveling Salesman problem.

Because we do not aspire to solve this problem a simpler method of quantifying performance was developed. A "performance index" is introduced that compares foraging cost to foraging revenues. This index is defined as the ratio of the distance covered by a paddler in a given period (D_t) and the product of the number (G_e) of glowballs eaten during that period and the average nearest neighbor distance while consuming a glowball ($\langle D_m \rangle$). If the paddlers on average move to the nearest neighbor after eating a glowball the performance index will be about 1. When no further information on glowballs or routing is available, this is the optimal foraging strategy [24]. Lower values indicate lower performance; the possible significance of higher values is discussed below.

$$\wp = \frac{G_e \cdot \langle D_m \rangle}{D_t} \tag{1}$$

The average distance between nearest neighbors $\langle D_m \rangle$, as used in calculating the performance index, is calculated by an algorithm that starts with the glowball nearest to the center of the foraging environment; determines the distance to its nearest neighbor; finds the nearest neighbor of that glowball, and so on. Hence the distances measured reflect the minimal distance to the nearest uneaten glowball. As a result, some distances may be quite large! Values of the index significantly larger than 1 are an indication that the index has become invalid, usually due to an overestimated $\langle D_m \rangle$ value.

3.2 The Simulated Environment

Three paddlers were allowed to roam an unlimited patch of deep sea. An open, rectangular section of this space served as a foraging environment: In this foraging space, a number of glowballs were distributed. At fixed intervals, eaten glowballs were replaced (i.e., made visible to all paddlers) at their initial position, to prevent depletion.

An experiment consists of an observation of the three identical paddlers over a certain time, during which data are collected, averaged over the three paddlers, and stored at the end of the observation time. Paddlers venturing too far from (i.e., losing visual contact with) the foraging space are replaced at a random position within the foraging space. This procedure was repeated for each value out of a fixed range of the

parameter of interest. During these experiments, there is no interaction or competition between the different paddlers: They do not exist for each other, and a glowball eaten by one remains visible and edible for the other two. Thus we collect data on three independent,[3] identical paddlers in the same environment and with the same settings in parallel. This is fast and convenient in our setup but has no principled significance.

3.3 Simulation Parameters

Experiments were carried out using a proprietary simulation package written in ANSI C and run on HP 9000/730 and Apollo DN10000 computers. The simulated two-light experiments on the archaepaddlers L and B are described in Section 4. In all other experiments, B-type paddlers were "released" into the foraging space, which was 150 length units square. They were observed for 4,000 or 5,000 steps (time units), each with a resolution of $\delta t = 0.033$ (i.e., 30 clock ticks per time unit of $t = 1$), during which data were collected. Glowballs numbered between $G = 64$ and $G = 625$ distributed uniformly; eaten glowballs were replaced every 500 or 1,000 steps. The glowballs had a radius of 0.5 ± 0.144 length units (uniform distribution; *large individuals*). In some experiments, *small individuals* of radius 0.02 ± 0.00289 (also uniform) were used. Glowball luminance was always set at 50 ± 0.279 (uniform). Glowballs closer by occlude those that are farther away. The paddlers had a radius of 2 length units and a mass of 500 mass units. We use the values listed in Table 1 as default values for the paddler's different parameters; they are part of a reasonable and reasonably flat optimum in the fitness landscape ($\wp$ as a function of the parameters).[4]

Experiments were done in two types of environment. We will first discuss the experiments performed in an environment with absolutely clear water (without any absorption of light), and in which an ad hoc, relatively "close," hard visual horizon (D) is imposed. Then we will discuss the effects of a more realistic environment in which a substantial amount of light is absorbed by the water, the glowballs are less bright, and the paddlers' thresholds higher. In these experiments, the visual horizon is a result of available light, absorption, and internal thresholds; these are tuned such that the paddlers can see glowballs of average brightness up to a distance $\approx D$.

We did not investigate the possible influence of noise in the paddlers' nervous systems. The paddlers consist of analog components, and the change of parameter values never leads to catastrophic changes of performance. Such changes always proved to be gradual. Thus noise cannot have catastrophic effects either. Moreover, substantial spatiotemporal averaging is performed during the simulations. Therefore we are confident that performance will merely degrade gracefully with noise.

4 Experiments and Results

4.1 Experiment 1: The Two-Light Experiment

To assess to what extent our hypothetical animal's navigation might be representative of a simple, real animal's behavior, we performed some two-light experiments [12]. In the more realistic, absorbing environment, which will be described in full detail below (see also the entries for the absorbing medium in Table 1), two identical, "large" glowballs were placed at a distance of 30 length units from each other (size 0.5, luminance 0.0005). A single paddler was repeatedly placed approximately on the line through the midpoint of, and perpendicular to, the line segment connecting the two glowballs. The initial position and orientation were slightly varied across trials (of course, exactly identical

3 They do not "exist" for each other, meaning they can even occupy the same physical location, nor is there interaction through competition.

4 Actually, most combinations of these parameters gave rise to some foraging success, except of course for pathological cases (like eyes and paddles at the same position).

Table 1. Default parameters of the archaepaddlers.

parameter	*value*	*comment*
α	90°	eye position; varied in Experiment 2
β	20°	eye axis; varied in Experiment 2
φ	140°	width of the eyes' field of view
σ	15°	halfwidth of the RWF (B-type paddlers) varied in Experiment 3
	35°	for the absorbing medium
D_f	11.5	fixation distance in length units varied in Experiment 4
ε	-30.0154°	the visual axis: the value listed corresponds to $D_f = 11.5$
	-23.1847°	($D_f = 36$) for the absorbing medium
D	10	visual horizon in length units
	25	for the absorbing medium
κ	45°	angular width of the mouth length of the mouth: 1.46 length units
δ	130°	paddle position
ρ	-20°	paddle rotation
ϖ	0.5	binocular weight (B-type paddlers)
scotopic threshold	10^{-10}	lowest luminance level at which vision can be used
	10^{-5}	for the absorbing medium

initial conditions will give rise to identical visually guided paths). The paths taken by the paddler were recorded and are shown in Figure 4. For comparison, Figure 4a shows the paths for *Corophium longicorne* (Figure 73, [12]). Figure 4c shows the paths taken by a B-type paddler as described above. Figure 4b presents the paths taken by an L-type paddler with the same parameters (but of course without directional selectivity ($\sigma = \infty$) and no extra binocular weighting). The B- and L-type paddlers were released from the same starting positions (cf. the legends in Figures 4b, c).

There is a striking difference between the paths taken by the L-type (Figure 4b) and the paths taken by the B-type (Figure 4c) paddlers. The L-type paddler's behavior resembles much more the behavior of the real animal (Figure 4a) than does the behavior of the B-type paddler.[5] It can clearly be seen that the RWF and the extra binocular weighting significantly increase the paddler's apparent "decisiveness" (Figure 4c). From these data, one can conclude that *Corophium l.* has few or no mechanisms to increase its directional selectivity. The L-type paddler is a satisfactory model. The following additional observations can be made. Each point of the traces in Figures 4b and 4c represents the center of the paddler's body; the switch to searching behavior after having eaten a glowball thus seems to occur at some distance from the glowball. The amplified continuation of the last maneuver is noticeable. Here the distance between the two glowballs is such that initially both glowballs are visible, but at a certain moment, the

5 One might say that the L-paddler seems to show more directional selectivity than *Corophium l.*, albeit that it has no directionally selective mechanisms. This must be a result of different characteristics of the visual systems, and possibly of the environment.

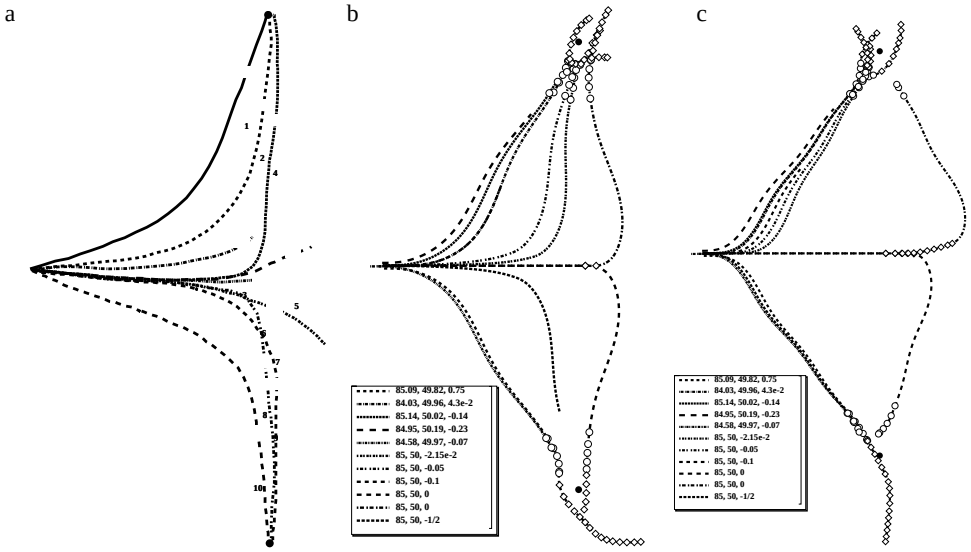

Figure 4. a) Paths followed in a two-light experiment by a marine amphipod, *Corophium longicorne*, in two equal lights. After [12]. b) Paths followed in a two-light experiment by a paddler without directional selectivity, presented with two identical glowballs. Closed circles represent the glowballs (at (100,65) and (100,35); open circles mark the route segments where the paddler was eating (had the glowball in its mouth); diamonds mark the route segments where the paddler was searching. The legend lists the starting coordinates (x, y, φ), where (x, y) are the starting position and φ the orientation with respect to the line of equal excitation (the line midway between the two glowballs). c) Paths followed in a two-light experiment by a paddler with directional selectivity, presented with two identical glowballs.

paddler's movement will cause one of the glowballs to disappear out of one eye's sight.[6] This explains why the paddler in most cases "chose" one glowball or the other—which it continues to do until the inter-glowball distance becomes substantially smaller than D. Lastly, even when the paddler is initially trapped in a donkey's dilemma (i.e., when it receives equal stimulation on the left and right), its searching mechanism causes it to choose one glowball (at random) after having passed both. As the donkey's dilemma is an unstable equilibrium, it is most unlikely that a paddler or indeed a real animal will ever get into one during its "real life." We will now turn to studying the B-type paddler.

4.2 Experiment 2: The Influence of Eye-Position Parameters α and β

In this experiment, we study the influence of eye position on foraging performance. To this end, the eye's position on the body was systematically varied from fully frontal to fully lateral (i.e., $\alpha \in [0^\circ, 90^\circ]$). The direction of the eye axis was varied independently, such that the eyes were either looking both outward, or both inward ("squint"), and a number of directions in between (i.e., $\beta \in [-90^\circ, 90^\circ]$). The other parameters were as listed in Section 3. D_f was kept at the value listed in Table 1 (thus ε covaries with β), except for fully frontal eyes; in that case, the eyes' visual axis is along the eye axis ($\varepsilon = 0^\circ$ when $\alpha = 0^\circ$ and we define $D_f = 0$). The results are shown in Figure 5 for a low glowball density and in Figure 6 for a high glowball density, in both cases for populations of large glowballs.

6 The geometry of the paddler's field of view is such that when following the line midway between the glowballs, the glowballs are lost out of sight when the paddler's center passes the glowballs.

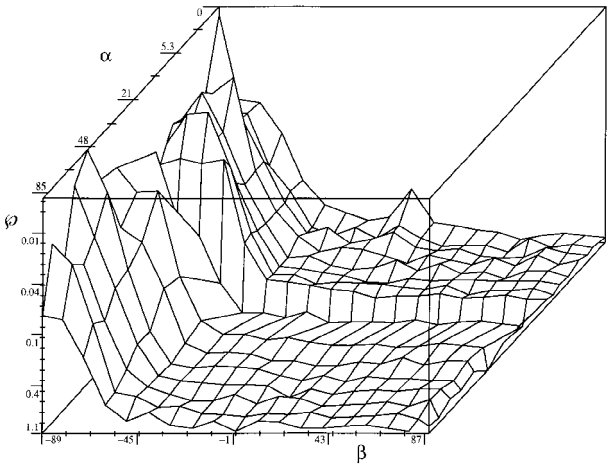

Figure 5. Influence of eye-position α and eye-axis β on the performance for a low density, $G = 121$, of large glowballs. For clarity, the $\wp$ axis is reversed.

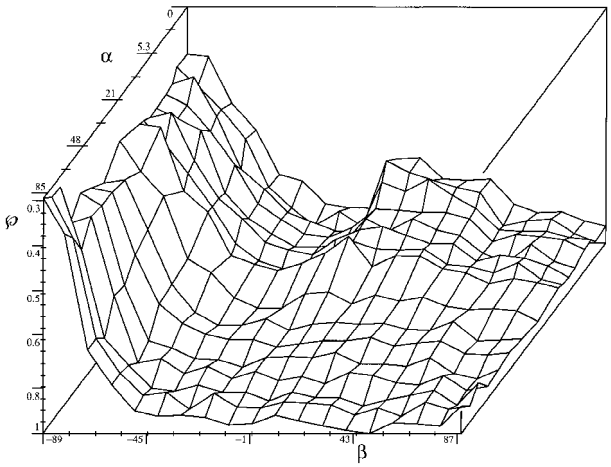

Figure 6. Influence of eye-position α and eye-axis β on the performance for a high density, $G = 625$, of large glowballs.

In both figures, it can be clearly seen that there is a very broad range of eye position/eye axis combinations that allow adequate foraging. Contrary to what might be expected intuitively, this is especially true for *low* glowball densities (Figure 5). At high glowball densities (Figure 6), eye position/eye axis tuning is sharper, in the sense that some combinations give rise to highly inferior performance. For most suboptimal combinations (generally $16 \leq \beta \leq 64$), performance is better than for low glowball densities, possibly due to the increased probability in high densities of blindly stumbling upon a glowball. Optimal performance is the same in both densities. For low glowball densities of $G = 121$, the average distance between nearest neighbors has at least the value of the visual horizon ($\langle D_m \rangle \geq D$), meaning that there is on average just one glowball visible at any given moment. As densities are increased, the number of glowballs visible at any time also increases. As a result, equalizing the luminance responses in both eyes does not necessarily lead the paddler to a glowball. It may even

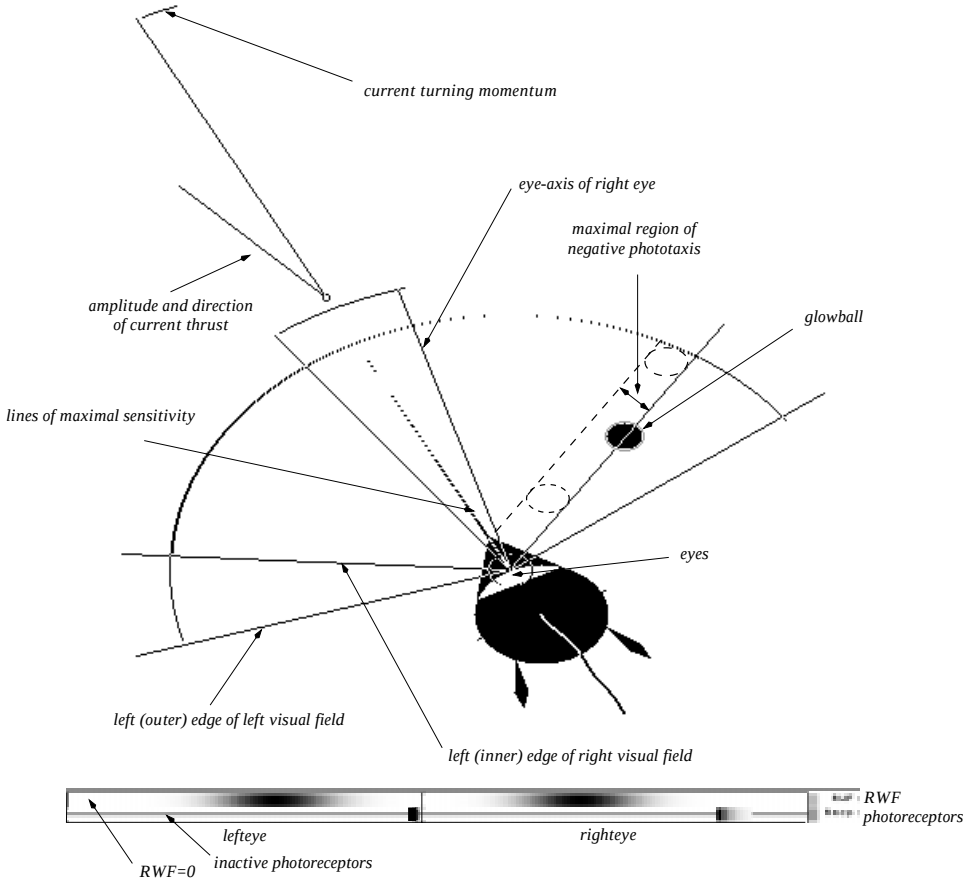

Figure 7. A case of negative phototaxis. A screenshot of a paddler with $\alpha = 1^\circ$, $\beta = 10^\circ$, and $D_f = 11.5$ (hence $\varepsilon \approx -10^\circ$). The small lower panel shows the retinal weighting function (RWF) and the activity in the photoreceptors, color-encoded with black equal to maximal activity, and white equal to 0 activity or weight.

cause it to take a middle course between two equally attractive glowballs and thus miss both. Of course, such indecisiveness reduces the foraging effectiveness; thus the eye position/eye axis tuning becomes sharper.

In the case of extreme squinting (strongly negative values of β), and/or fully frontal, squinting eyes ($\alpha = 0^\circ$, $\beta < 0^\circ$), performance decreases dramatically. This is the result of *negative phototaxis* because now each eye is looking at the wrong (contralateral instead of ipsilateral) half of the world. As a result, paddlers turn rapidly away from every glowball entering their visual field. This explanation is supported by the fact that for eyes placed in this manner, the paddlers have on average smaller eye responses (i.e., they shun light, which minimizes retinal illumination) than paddlers with normal positive phototaxis.

When there are many glowballs, frontally placed, outward-looking eyes (low α and low, positive β, respectively) cause a drop in performance. This is the result of an interesting phenomenon, which arises from the way the eye's response is determined. Along the inner edges of the binocular visual field (the part sampled by both eyes), phototaxis turns out to be negative. This is because the eye response is taken to be the sum of all photoreceptor responses, divided by the number of supraliminally activated

receptors. See, for example, Figure 7. The paddler is approaching a glowball that was initially visible only to the right eye and has now entered the left eye's visual field for approximately 50%. The "vanes" in front of the paddler show its current movement (momentum) and the thrust currently generated. It can be seen that the paddler tries to swim quite vigorously in a direction opposite to that one would expect: negative phototaxis. The explanation can be gleaned from the "retinal panel": in the left eye, a few photoreceptors are stimulated, but all more or less to the same amount. In the right eye, the sum of the photoreceptor responses is higher, but due to the larger range of responses and the larger number of active receptors, the eye's response turns out to be lower than the left eye's response! This phenomenon depends on RWF and on glowball size: As indicated in Figure 7, the part of the visual field where it occurs is maximally as wide as the glowball inducing the phenomenon. It should be noted that the paddler does not always turn away from the glowball as a result: When its initial momentum is high enough, it will still be able to turn successfully toward the glowball. In this case, it happens that in the right eye the sum of the photoreceptor responses is higher than in the left eye, but its response is actually smaller than the response of the left eye. In the rest of the overall visual field (the combined visual fields of both eyes; surrounding the negatively phototaxic parts), phototaxis is positive. In the middle of a glowball population, behavior of such a paddler is quite normal. However, a dense population contained within a restricted area possesses relatively distinct borders. Once a paddler with this placement of the eyes reaches such a border, it will in many cases show a qualitatively different behavior, following the border—clearly no longer foraging. This edge-following behavior results from the fact that the glowball population as seen by the outer, positively phototaxic part of a visual field will cause the paddler to turn toward the population. Due to the high density, chances are very high that the small negatively phototaxic part of the eye's visual field will receive stimulation—in fact, this occurs whenever a glowball "moves" from the monocular part into the binocular part. The paddler can respond quite vigorously to such an event, as can be seen from Figure 7. As a result, the paddler turns away again, losing the negative stimulation. When the glowballs on the edge are close enough together for the phenomenon to repeat, the result is a slightly zigzagging, edge-following movement.

At medium to low glowball concentrations (only results obtained in low concentrations are presented here), chances are low to very low that the negatively phototaxic part of the visual field receives much stimulation because on average the distance between glowballs is larger than the visual horizon ($\langle D_m \rangle > D$). Also, such a population does not show distinct borders. This explains why edge following is only seen at high glowball concentrations.

The results of Figures 5 and 6 show that the default values for eye position and eye axis ($\alpha = 90^\circ$ and $\beta = 20^\circ$, respectively, as listed in Table 1) are compatible with a high performance. The orientation of the eye axis (β) is the parameter most dramatically affecting foraging performance in these environments.

4.3 Experiment 3: The Retinal Weighting Function

In the third set of experiments, we address the influence of the halfwidth σ of the eye's retinal weighting function on the foraging performance. It might be expected that intermediate values give the best performance: They promote selectivity (especially at high glowball densities) while not sacrificing too much of the visual field. Figures 8 and 9 show the results of two simulations: one with large glowballs, and one with small glowballs, in both cases for four different target densities from low (64) to high (625).

For low glowball density and the narrowest halfwidth, the paddlers are almost functionally blind and thus spend most of their time in the searching mode. The performance index turns out (conveniently) to be approximately 0.5 in that case.

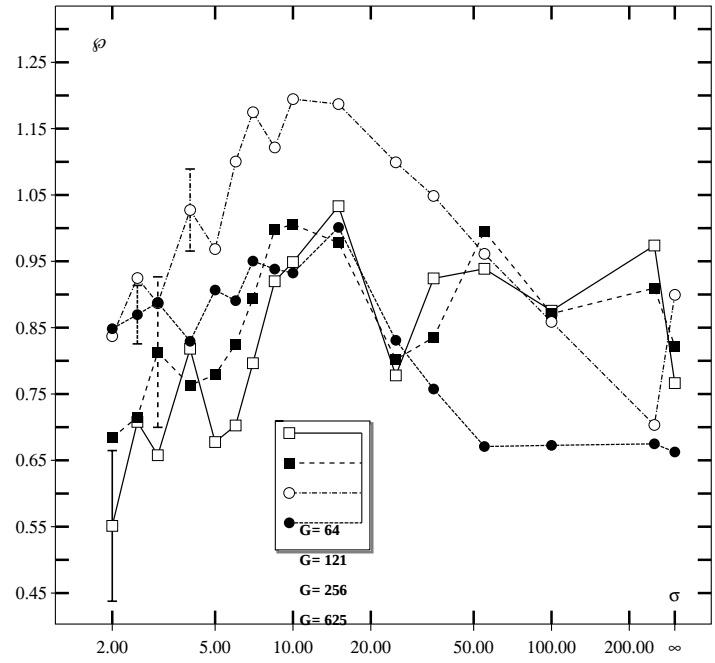

Figure 8. Influence of RWF halfwidth σ on the performance for various glowball populations with large individuals. Error bars indicate the average standard deviation (measured over the performance of three paddlers) per curve.

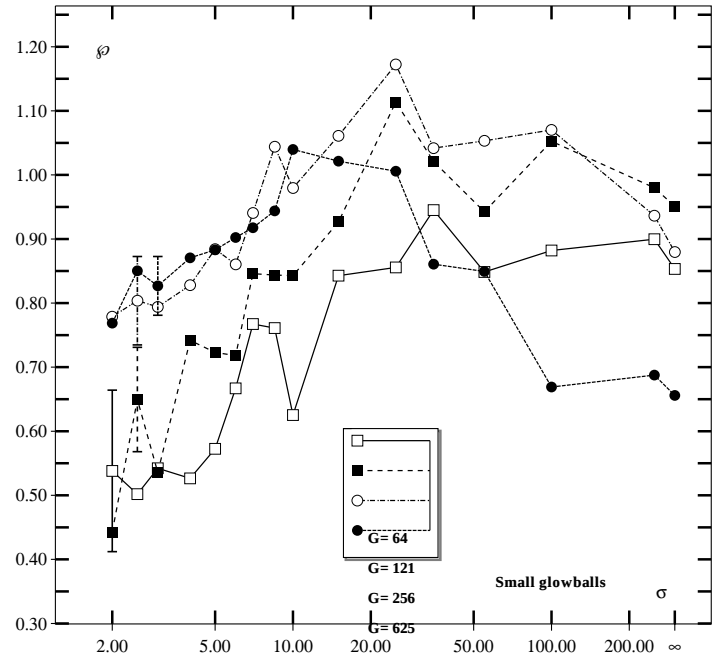

Figure 9. Influence of RWF halfwidth σ on the performance for various glowball populations with small individuals. Error bars indicate the average standard deviation per curve.

Intermediate halfwidths indeed give the best performance, but only for higher glowball densities, and most clearly for large glowballs. Roughly speaking, we find that the halfwidth should lie between 10° and 50°, that is $3 \cdot \sigma < \varphi$, meaning that the Gaussian distribution is narrower than the visual field.[7] When there are only few glowballs, a larger halfwidth (meaning a larger *functional* visual field) is no longer a disadvantage: There are simply no glowballs to be "distracted by."

In populations with small glowballs, the overall picture remains the same, except that performance does not degrade that much for larger halfwidths. Because small glowballs are easier to miss (due to limited retinal resolution and sensitivity, and to a lesser extent their smaller size, which demands higher navigating precision), a larger functional visual field increases foraging success, and the best performance occurs at somewhat larger halfwidths (except again for high densities—$G = 625$).

Performance is best at low-to-intermediate glowball densities, where the average number of visible glowballs is approximately one. This is of course due to the fact that the present simple positive phototaxis strategy does not allow for selective attention. When more than one glowball is visible at a given moment, the paddler tends to head toward the "center of gravity" of the luminance distribution that these glowballs project on its retinas. Under these circumstances, it might then miss most (if not all) of the targets.

4.4 Experiment 4: The Influence of the Angular Distance Between Eye Axis and Visual Axis

This experiment addresses the influence of the direction of maximal sensitivity of the retinal weighting function, the visual axis, as expressed in terms of the fixation distance D_f. Figures 10 and 11 show results for moderately low ($G = 121$) and high ($G = 625$) glowball densities, and a number of eye positions (α). To get a stronger dependence of performance on the fixation distance, the halfwidth of the retinal weighting function has been reduced to $\sigma = 5°$ in these experiments.

The results show first of all that the fixation distance should not be (much) smaller than the visual horizon. Again, the reason is the phenomenon of negative phototaxis, which occurs when the eyes are most sensitive to targets on the wrong (contralateral) side of the body. An additional factor at small fixation distances is that an increasing part of the visual field becomes unreceptive to visual stimulation. In other words, the paddler becomes increasingly visually handicapped (one might interpret it as tunnel vision).

It can also be seen that for fixation distances above 10 length units, performance remains more or less optimal. There is some indication that performance is best for fixation distances between 10 and 15 length units, that is, with the binocular point of highest sensitivity at or just beyond the visual horizon (D). This also holds for visual horizons $D \in \{5, 25, 50\}$ length units, indicating that the direction of optimal sensitivity is closely coupled to the visual horizon (a property of the medium).

4.5 Experiment 5: The Influence of a More Realistic, Absorbing Medium

The fact that the optimal visual axis is related to the visual horizon raises the question of what it would be for a more realistic description of absorption in the medium. We therefore repeated the above experiments after including a model of light absorption by water. Luminance L as a function of distance (d) to the source (with luminance L_0) in a medium with an absorption coefficient Ω, is given by (see, e.g., [23])

$$L[d] = L_0 \cdot e^{-d \cdot \Omega} \tag{2}$$

7 That is, depending on ε, the visual field contains approximately 95% of the distribution.

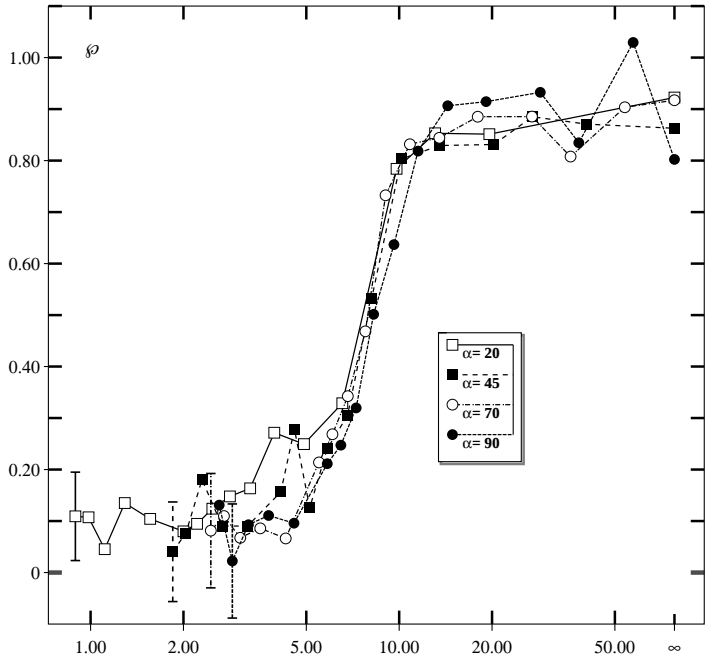

Figure 10. Influence of fixation distance D_f on the performance for $\sigma = 5^\circ$, various eye positions α and a low glowball density of $G = 121$. Error bars indicate the average standard deviation per curve.

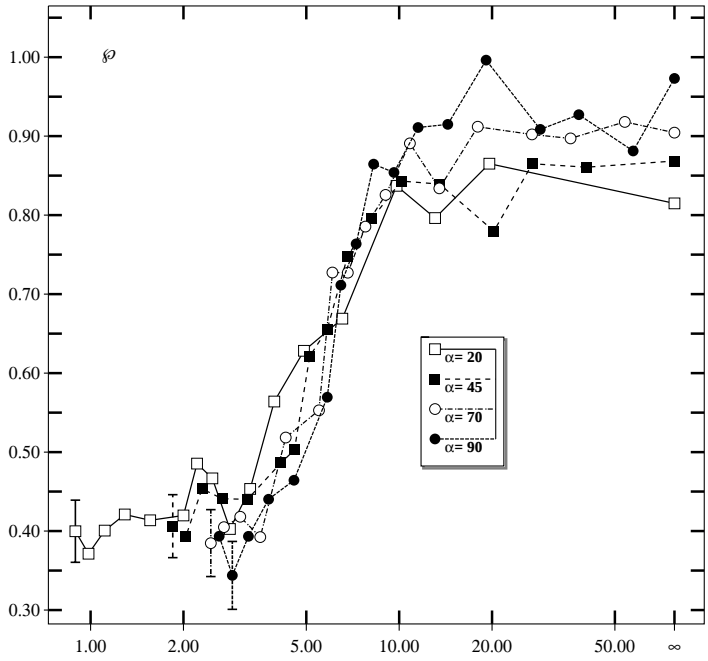

Figure 11. Influence of fixation distance D_f on the performance for $\sigma = 5^\circ$, various eye positions α and a high glowball density of $G = 625$. Error bars indicate the average standard deviation per curve.

The threshold for switching to searching behavior and the threshold of the photoreceptors were set to a higher value (10^{-5}). Glowball luminance was reduced to $5 \times 10^{-4} \pm 3.14 \times 10^{-5}$. Absorption by the medium was chosen in such a way that a glowball of average luminance remained visible up to a distance of 25 units (thus $\Omega \approx 0.156$; the visual horizon is now no longer a "hard" boundary).

A striking effect of such a strong absorption is a phenomenon that might be described as "binocular absorption parallax." Due to unequal target-to-eye distances, the two eyes can receive different illumination from a single target, even in cases where the target excites the same number of photoreceptors in both eyes. This can lead to zigzag locomotion for body sizes that are large in relation to the visual horizon as set by absorption.

In general, however, the results were qualitatively similar to the results described above. Absorption by the medium acts as a filter comparable to the retinal weighting function, in the sense that it filters out more distant targets. Therefore, a broader halfwidth ($\sigma \approx 35^\circ$) is favored (this increases the probability of detecting a glowball). The optimal fixation distance (for $\sigma = 5^\circ$) is slightly larger than it would be with a hard visual horizon at the same distance: $D_f \approx 36$. This reflects the fact that the absorption horizon is not a clearly defined border: Brighter-than-average glowballs can be seen over a larger distance than less bright glowballs. No effect was found on the optimal eye-position parameters.

5 Discussion

5.1 Phototaxic Navigation: A Robust Strategy

Loeb's (conceptual) model describes phototaxic foraging in the two-light experiments very well for many simple species, but as shown by our results, it also works quite well in an environment with many targets. Two innovations improving course stability were added: a retinal weighting function, which gives directional selectivity, and extra weighting of binocularly visible targets, which serves as a crude fixation mechanism. We then performed a number of simulations to study how habitat parameters and body geometry influence performance. We interpret the results of those simulations as an indication that a compromise should be found between, on the one hand, the ability to detect and react to as many targets as possible, and on the other hand, the ability to concentrate on one or a few attractive targets (in terms of size, direction, and proximity). The Gaussian directional selectivity shows this clearly for higher glowball densities (see, e.g., Figure 8): When the selectivity curve is too narrow, performance is poor, because only the nearest or brightest glowballs are seen; when it is too wide, performance is poor because the paddler cannot decide which of all visible glowballs to select. Using classifier systems to evolve a reimplementation of Wilson's animat, Cliff and Bullock [7] reported similar results: The optimal pattern of sensory sampling ("foveal vision") depends on the animat's environment.

Within these constraints, a surprisingly large freedom of parameter choice exists. Adequate foraging performance results for many, even intuitively disadvantageous parameter settings. This might be taken to indicate that our quantification of performance in terms of the index $\wp$ is unsatisfactory. We think this is unlikely because it compares performance against the strategy "go to the nearest visible neighbor," which is the optimal strategy in this paradigm [24].[8] It therefore appears more likely that a wide range of parameter variation is acceptable in an "easy" environment that does not pose strict requirements (on foraging strategies). In other words, this finding suggests that the overall design of this type of autonomous agent is very robust.

8 And indeed performance indices measuring energy expenditure give similar results.

Being able to detect as many glowballs as possible can be accomplished by having the eyes as widely apart as possible, and "looking" slightly outward (a larger visual field). Indeed we do find a slight performance increase for such a geometry. Eyes looking too much inward and set too close together can cause parts of the visual field to become negatively phototaxic: In those regions the (squinting) eyes respond better to contralateral then to ipsilateral targets. This can result in a substantial performance decrease. Such a geometry can even result in a visuomotor system that follows the edges of (dense) glowball populations in an endless loop of attraction/repulsion by glowball light. An extreme convergent squint decreases the paddler's fitness.

From these results we derive two predictions for real animals using a similar simple phototaxic foraging strategy. In the first place, we predict that the within-species variation of the eye position can be quite substantial. But, in the second place, we predict that squinting will be virtually absent in the populations. We are not aware of quantitative studies in real animals that allow us to test these predictions at present. On the whole, there is an amazing paucity of statistical data on the geometric parameters of animal bodies. Our study shows that such data would be important in connection with theories on, or models of, navigation systems.

Our earlier experiments on extremely light paddlers (that do not need a speed control system because acceleration and deceleration are almost instantaneous) favored intermediate eye positions that resulted in more stable locomotion (less "wobbly"). Locomotion that is too wavery tends to cause more "glowball misses," thus impairing foraging performance. The more biologically plausible heavier paddlers discussed in this article do need the speed control system; due to the larger inertia, locomotion then even remains stable for extreme eye positions. This leads to the prediction that there is a positive correlation between mass and the presence and quality of speed control.

It must be noted that the robustness we find applies only to 2D movements, in the plane perpendicular to the animal's plane (or line, in our 2D animal) of bilateral symmetry. Using the current sensorimotor model, a 3D animal with 2D retinas will show the same robustness when it is only concerned with, say, horizontal locomotion. We predict that similar robustness can be obtained for movements in 3D when the animal is equipped with a locomotor system that can generate left and right vertical thrust, by simply dividing the retinas in an upper and lower part, and projecting, for example, the left upper part to the right upward motor.

5.2 The Influence of the Eyes' Directional Tuning

A similar reasoning can be applied to the directional selectivity as specified by the retinal weighting function. Too narrow an RWF will reduce a (large) visual field to a substantially smaller functional visual field. Too wide an RWF will lead to unstable locomotion: Peripheral glowballs can then cause the paddler to wander away from a close glowball straight-ahead. Simulations in which the width of the RWF was varied support this: Performance is best for an intermediate width. In low-density glowball populations there is less risk for distraction by peripheral glowballs so performance does not degrade as much for wider RWFs.

The direction of the maximal sensitivity of the RWF (ε, also specified as the fixation distance D_f) also has a substantial influence; performance is best when $D_f \approx D$. More inward tuning ($D_f < D$) causes parts of the visual field to become negatively phototaxic. More outward tuning ($D_f > D$) places too much sensitivity in the periphery. In addition, in this case sensitivity in the area in front of the paddler changes rapidly with direction (being "covered" by the flanks of the Gaussians), meaning that differences in eye response to different targets in that region can become too large. All these effects cause unstable locomotion or too little "concentration." The constraints on the compromise between directional selectivity and fixation capability increase with increasing

glowball densities; all studied paddlers have difficulties with successful foraging when several glowballs are visible at any given moment. This problem can only be solved by the evolution of a selection mechanism that allows a paddler to fixate (concentrate) on a single glowball, an interesting challenge that might lead to a more advanced species of paddler.

5.3 Advantages of the Tinkering Method

The present model has been specified in such a way that it is easy to evolve. In principle, this could be done using genetic algorithms [16] or classifier systems [17], which simulate evolution from randomly disturbed inheritance of parameters and a competition for food among members of the generations that coexist. However, this is a very time-consuming technique, both computationally and from an implementation point of view. We would have to solve the problem of how to encode a nervous system in a genome, which is not directly relevant to our interests. A growing number of studies report on just such approaches (e.g., [3, 8]), including research shedding evolutionary light on the relation between body form and behavior that we find [26, 27]. Also, it turns out that phototaxic, paddling systems are an evolutionary possibility [26]. It would be very interesting to see what Sims' model would come up with when neurons with a more limited and biologically acceptable set of "IO functions" were used. However, as outlined by Sims, such techniques allow the creation of systems "without requiring an understanding of the procedures or parameters" (p. 15), or indeed "it can be difficult to analyse exactly how a control system [. . .] works" (p. 18). Even if such an approach were far less problematic, our method of manual evolution, that is, "designing" animal classes and refining them with parametric simulations and by adding, changing, or deleting connections, neurons, and so forth, might still be preferable in many cases. It provides a deeper insight in the influence and interactions of the various parameters of the animal-habitat coupling mechanisms. In short, tinkering [19] at the organismic and network level can be a powerful method of studying the possible adaptations (design space) of hypothetical and real animals in quantitative detail. This approach is not inferior to evolutionary simulations if it is constrained by biological knowledge and combined with parametric simulations that can show us what the stable and optimal regions are in design space.

5.4 Future Work

The paddlers provide a convenient platform to develop and test more complex foraging strategies and/or to test hypotheses about visual processing or motor control in real animals. Some examples of future research opportunities, directly inspired by our current findings, are

- To simulate neural networks that can internally represent direction, distance, or size of all the glowballs in sight (or even those remembered), and to use this "knowledge" to optimize the choice of the next target.
- To introduce maplike memories of regional visual analyses so that the paddler can plan a path before it "hits the road."
- As our insight into the physics of locomotion deepens, to improve on the present simple specification, and to test more sophisticated theories of neuronal control of locomotion.

We have already reported on mechanisms for dark/light adaptation in the diurnal paddler, a close relative of the B-type archaepaddler [5]. Research is currently under

way to evolve a motion detection system from the dark/light adaptive system, and to use the new motion parallax information to improve visual analysis of relative distances and/or relative motion of external objects.

Acknowledgments

The research presented in this article was part of a Ph.D. project supported by a grant from the Foundation for Biophysics (presently Foundation for the Life Sciences) of the Netherlands Organisation for the Advancement of Scientific Research (NWO). The first author wishes to thank the following people who at crucial moments lent their insights to increase the article's "fitness" during its evolution: Professor P. Hogeweg of the Bioinformatics Department at Utrecht University, and Dr. J. Droulez and R. Grasso at the LPPA, College de France in Paris.

Appendix: Locomotion

An overview of the geometry of locomotion is given in Figure 12. When $\rho \neq 0^\circ$ the thrust vector generated by both paddles has a rotational component and a translational component. Because the rotational component of the thrust vector is generally small with respect to the translational component, and the paddles themselves generate only a very small drag, a stiff tail fin has evolved to increase rotational drag. Due to this tail, drag-induced torque keeps the paddler's long axis aligned with the direction of movement.

We assume that the paddler can be considered a solid body moving in a homogenous, laminar medium. Thus simple mechanics (*Stokes' law*; see, e.g., [36]) tells us that its translational speed (v) can be approximated using the following equations:

$$v^{\bullet} = \frac{T_{\text{trans}} - 0.5 \cdot D \cdot f_d \cdot v^2}{M} \tag{3}$$

T_{trans}, the resultant translational thrust, can be expressed as a constant ($f_s = 1$; the speed factor) times the translational component of the thrust vector (F_{trans}):

$$T_{\text{trans}} = f_s \cdot F_{\text{trans}} \tag{4}$$

The resultant frontal drag factor, f_d, depends on the frontal drag coefficient (C_w) of the paddler's body, and the current angle of movement (Δ):

$$f_d = C_w \cdot \left|\sin[\Delta]\right| + C_w \cdot \left|\cos[\Delta]\right| \tag{5}$$

The angular speed (ω) can be approximated with

$$\omega^{\bullet} = \frac{T_{\text{rot}} - 0.5 \cdot D \cdot C_{rw} \cdot \omega^2}{M} \tag{6}$$

T_{rot}, the resultant rotational thrust, depends on the rotational component of the thrust vector (F_{rot}) and the current speed and direction of movement:

$$T_{rot} = f_s \cdot \left(F_{rot} + C_w \cdot \left(v \cdot \sin[\Delta]\right)^2\right) \tag{7}$$

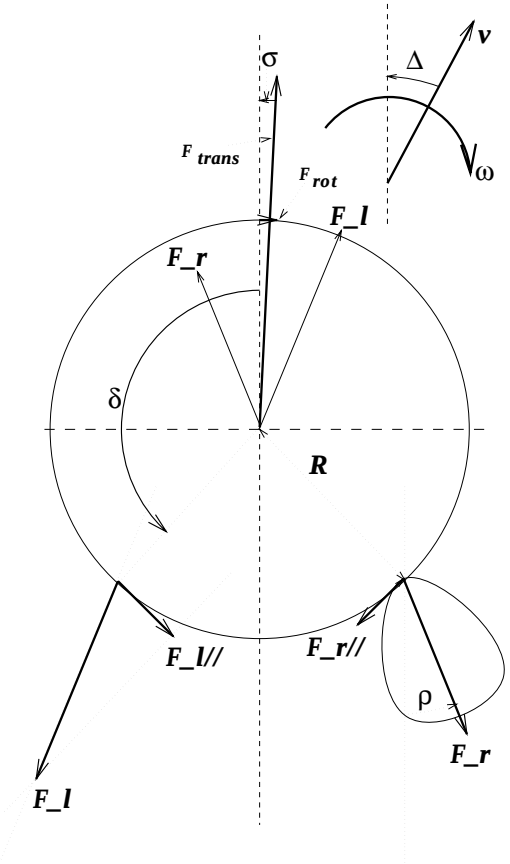

Figure 12. Overview of the forces exerted by a pair of paddles on the paddler body. F_{rot} is the resultant rotational component of the thrust; F_{trans} is the translational component, which works under an angle of σ.

with

$v^{\bullet}$ ($\omega^{\bullet}$) the current (angular) acceleration.

M the paddler's mass. We take $M = 500$.

D the density of the medium. We take $D = 1$.

The area of the tail fin is expressed in terms of the radius R of the paddler's body. The tail fin can be divided into two sections: one "above" the body and one extending behind the body. The part above the body has length R and height $\frac{1}{4} \cdot R$, which is half the body height. The part extending beyond the body has length $\frac{1}{2} \cdot R$ and height $\frac{3}{4} \cdot R$. Paddle-induced drag is neglected, which gives us a radial drag coefficient $C_{rw} = \frac{65}{96} R^4$.

References

1. Armstrong, D. M. (1988). The supraspinal control of mammalian locomotion. *Journal of Physiology* (London), *405*, 1–37.

2. Beer, R. D. (1990). *Intelligence as adaptive behaviour—An experiment in computational neuroethology*. Boston: Academic Press.

3. Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior, 1*, 91–122.

4. Bertin, R. J. V. (1994). *Natural smartness in hypothetical animals—Of paddlers and glowballs*. Unpublished doctoral dissertation, Utrecht University.

5. Bertin, R. J. V., & Grind, W. A. van de. (1997). The influence of light/dark adaptation and lateral inhibition on phototaxic foraging. A hypothetical-animal study. *Adaptive Behavior, 5*, 141–167.

6. Braitenberg, V. (1984). *Vehicles—Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

7. Cliff, D., & Bullock, S. (1993). Adding "foveal vision" to Wilson's animat. *Adaptive Behavior, 2*, 49–72.

8. Cliff, D., Husbands, P., & Harvey, I. (1993). Explorations in evolutionary robotics. *Adaptive Behavior, 2*, 73–110.

9. Corbacho, F. J., & Arbib, M. A. (1995). Learning to detour. *Adaptive Behavior, 3*, 419–468.

10. Cruse, H., Brunn, D. E., Bartling, Ch., Dean, J., Dreifert, M., Kindermann, J., & Schmitz, J. (1995). Walking: A complex behavior controlled by simple networks. *Adaptive Behavior, 3*, 385–418.

11. Ekeberg, O., Lansner, A., & Grillner, S. (1995). The neural control of fish swimming studied through numerical simulations. *Adaptive Behavior, 3*, 363–384.

12. Fraenkel, G. S., & Gunn, D. L. (1961). *The orientation of animals*. New York: Dover Publications.

13. Gibson, J. J. (1977). The theories of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing: Toward an ecological psychology*. Hillsdale, NJ: Erlbaum.

14. Grind, W. A. van de. (1990). Smart mechanisms for the visual evaluation and control of self-motion. In R. Warren & A. Wertheim (Eds.), *Perception and control of self-motion* (pp. 357–398). Hillsdale NJ: Erlbaum.

15. Hagevik, A., & McClellan, A. D. (1994). Coupling of spinal locomotor networks in larval lamprey revealed by receptor blockers for inhibitory amino acids: Neurophysiology and computer modelling. *Journal of Neurophysiology, 72*, 10–29.

16. Holland, J. H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, University of Michigan Press.

17. Holland, J. H., & Reitman, J. S. (1978). Cognitive systems based on adaptive algorithms. In D. A. Waterman & F. Hayes-Roth (Eds.), *Pattern-directed inference systems*. New York: Academic Press.

18. Holland, O., & Melhuish, C. (1996). Some adaptive movements of animats with single sensor symmetrical sensors. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 55–64). Cambridge, MA: MIT Press/Bradford Books.

19. Jacob, F. (1977). Evolution and tinkering. *Science, 196*, 1161–1166.

20. Kühn, A. (1919). *Die Orientierung der Tiere im Raum*. Jena, Germany: Gustav Fischer Verlag.

21. Locket, N. A. (1977). Adaptations to the deep-sea environment. In F. Crescitelli (Ed.), *Handbook of sensory physiology Vol. VII/5: The visual system in vertebrates* (pp. 67–192). Berlin/New York: Springer Verlag.

22. Loeb, J. (1918). *Forced movements, tropisms and animal conduct*. Philidelphia: Lippincott. (Republished 1973. New York: Dover)

23. Lythgoe, J. N. (1979). *The ecology of vision*. Oxford: Clarendon Press.

24. Rössler, O. E. (1974). Adequate locomotion strategies for an abstract organism in an abstract environment—A relational approach to brain function. In M. Conrad, W. Güttinger,

& M. Dal Cin (Eds.), *Lecture notes in biomathematics 4: Physics and mathematics of the nervous system* (pp. 342–370). Berlin/New York: Springer Verlag.

25. Schöne, H. (1984). Spatial orientation. The spatial control of behavior in animals and man. Princeton, NJ: Princeton University Press.
26. Sims, K. (1994). Evolving virtual creatures. In *Computer graphics, annual conference series* (pp. 15–23). New York: ACM SIGGRAPH.
27. Sims, K. (1994). Evolving 3D morphology and behavior by competition. *Artificial Life, 1*, 353–372.
28. Swain, M., & Ballard, D. (1991). Color indexing. *International Journal of Computer Vision, 7*, 11–32.
29. Terzopoulos, D., Tamer, R., & Grzeszczuk, R. (1996). Perception and learning in artificial animals. In C. G. Langton & K. Shimihura (Eds.), *Artificial life V: Proceedings of the Fifth International Conference on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press.
30. Terzopoulos, D., Tu, X., & Grzeszczuk, R. (1994). Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world. *Artificial Life, 1*, 327–351.
31. Trullier, O., Wiener, S. I., Berthoz, A., & Meyer, J.-A. (1997). Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology 1997, 51*, 483–544.
32. Walter, W. G. (1950). An imitation of life. *Scientific American, 182*, 42–45.
33. Walter, W. G. (1951). A machine that learns. *Scientific American, 185*, 60–63.
34. Wilson, D. M., & Waldron, I. (1968). Models for the generation of motor output patterns in flying locusts. *Proceedings of IEEE, 56*, 1058–1064.
35. Wilson, D. M., & Weis-Fogh, T. (1962). Patterned activity of co-ordinated motor units, studied in flying locusts. *Journal of Experimental Biology, 39*, 643–667.
36. Young, H. D. (1992). *University physics* (8th ed.). Reading, MA: Addison-Wesley.