The influence of light/dark adaptation and lateral inhibition on phototaxic foraging. A hypothetical-animal study

(7th March 2002)

Adaptive Behavior, Vol. 5, No. 2: 141 - 167

R. J. V. Bertin[1,2] (rbertin@ccr.jussieu.fr)
W. A. van de Grind[2] (W.A.vandeGrind@bio.uu.nl)

1: Current address:
LPPA
Collège de France / C.N.R.S.
11, place Marcelin Berthelot
75005 Paris
France
fax: +33 1 44271382
2: Neuroethology group
Department of Comparative Physiology
Utrecht University
&

Helmholtz Instituut
School for Autonomous Systems Research

Padualaan 8
3584 CH Utrecht
the Netherlands
fax: +31 30 2542219/2532837

The influence of light/dark adaptation and lateral inhibition on phototaxic foraging.

Abstract

Vision did not arise and evolve to just "see" things, but rather to act on and interact with the habitat. Thus it might be misleading to study vision without its natural coupling to vital action. Here we investigate this problem in a simulation study of the simplest kind of visually-guided foraging by a species of 2D hypothetical animal called the (diurnal) paddler. In a previous study, we developed a hypothetical animal called the archaepaddler, which used positive phototaxis to forage for autoluminescent prey in a totally dark environment (the deep-sea). Here we discuss possible visual mechanisms that allow (diurnal) paddlers to live in shallower water, foraging for light-reflecting prey in ambient light. The modification consists of two stages. In the first stage Weber adaptation compresses the retinal illumination into an acceptable range of neural firing frequencies. In the second stage highpass filtering with lateral inhibition separates background responses from foreground responses. We report on a number of parameter-studies conducted with the foraging diurnal paddler, in which the influence of dark/light adaptation and lateral inhibition on foreground/background segregation and foraging performance ("fitness") are quantified. It is shown that the paddler can survive adequately for a substantial range of parameters that compromises between discarding as much unwanted visual (background) information as possible, whilst retaining as much information on potential prey as possible. Parameter values that optimise purely visual performance like foreground/background segregation are not always optimal for foraging performance and vice versa. This shows that studies of vision might indeed require more serious consideration of the goals of vision and the ethogram of the studied organisms than has been customary.

Keywords: dark/light adaptation, lateral inhibition, hypothetical animal, phototaxic navigation, Weber adaptation.

1 Introduction

Vision is usually studied without direct recourse to suitable actions of the studied organism. Since the coupling of animals to their habitat concerns a loop of processes of the type "sensory → interpretative → motor → habitat → sensory → ...", purely visual studies might give a rather limited and biased view of visually-guided organism/habitat interactions (Gibson 1979; van de Grind 1990). In principle, electrophysiological studies of behaving animals provide one possibility to improve this situation. However, this approach is virtually impossible in the case of soft-bodied or small animals. Unfortunately it is in these animals that visually-guided behaviour should be studied in the first place to get some insight in its evolutionary beginnings and early development. Even in larger animals there is hardly a possibility to use "in-eco" electrophysiology for other than rather simple (non-vital) actions, such as making an eye-saccade or pushing a button. One alternative is offered by the present hypothetical-animal approach. In such a computer-simulation study one develops a hypothetical animal performing some natural action and then studies the influence of all processes involved in the continuous perception-action loop, while it is in place and complete. Of course, the initial approaches are bound to be a bit clumsy. One first needs an acceptable "platform", performing the simulated action with sufficient dexterity. The action has to be natural and it has to be simulated in a physically and biologically realistic way and in sufficient detail. Once a certain quality of the platform is reached one can "play" with all the interactions and — in the context of studies of vision — study the contribution of all sorts of visual mechanisms to the behaviour under study.

This general idea motivated us to embark on an extensive study of hypothetical animals that perform natural actions in a biologically plausible way, and in which the animal/habitat interactions are physically realistic. Our hypothetical animals do not mimic specific real animals; we aim for a realistic embodiment of specific perception-action loops. The success of this approach depends on the biological realism of each of the aspects of the perception-action loop and is not measured by the question whether or not the results reproduce the behaviour of a specific living animal. After all there are many missing links in the evolution of, say, navigation behaviour, but that does not preclude the development of reasonable theories as long as the known (physical and biological) boundary conditions are meticulously taken into consideration.

We have recently designed and simulated a species of hypothetical deep sea animal, called archae-paddler, that hunts autoluminescent prey, called glowballs, in an otherwise dark surround (Bertin & van de Grind 1996). Since the animal lives in the water-layer just above the bottom of the deep sea, we largely ignore the third dimension, essentially modelling the animal in 2D. Its foraging behaviour is phototaxic[1], and the rather primitive eyes drive the contralateral paddles.

There are several levels of detail in which acting animals can be modelled. In general, more complex behaviours are best modelled at a higher, more schematic level of description (e.g. Corbacho & Arbib, 1995), with less relevant, lower-level biophysical aspects approximated or taken for granted (e.g. the physics of locomotion; the animal might be able to move two steps forward and one to the side). However, since the animal's actions are its interactions with its environment, they directly influence its sensory input. Therefore studies in which sensorimotor behaviour is modelled at a lower level of description (e.g. Cliff, Husbands & Harvey 1993, Ekeberg, Lansner & Grillner 1995 or Cruse *et al.* 1995), and also studies aiming at realistic animation of behaviour (e.g. Terzopoulos, Tu & Grzeszczuk 1994) do explicitly model these physical aspects.

In our model, we simulate the physical processes of locomotion in sufficient detail, making the paddling realistic enough to view it as an adequate model of real-animal paddling (Bertin & van de Grind 1996). Since the anatomical and physical parameters proved to allow ample variation before the foraging quality began to break down, the simulated animal can be viewed as occupying a stable region of "design" space (in the sense of Dennett, 1995), allowing quite some genotypic variation. It is interesting to note that the geometry of such an animal depends on habitat-factors, such as prey-density and spatial prey distribution and on the simulated behavioural foraging strategy. Changes of the habitat and/or

[1] Kühn (1919), following Loeb (1918) defined *tropotaxis* as 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. This principle is nicely illustrated in Braitenberg (1984), but also in Walter (1950, 1951). Tropotaxis of chemical nature is demonstrated in Beer (1990).

behavioural strategy are immediately reflected in the body-geometry as required for optimal performance (Bertin and van de Grind 1996). Here we take the "average paddler" and "average habitat" of the previous study as our starting point and study the influence of increasing visual sophistication on the phototaxic foraging success. The foraging strategy, as embodied in the animal's nervous system, is the same as in the previous study: move in the direction of the strongest light response. This response is not an unprocessed reflection of the light distribution, since that would entail continuous course changes. Rather, the eyes have simple (Gaussian) tuning characteristics for direction, and binocular facilitation gives a slight bias for coursing straight ahead if there is prey in that direction. The visual information is thus preprocessed such that a balance is struck between on the one hand as large a field of view as possible (maximisation of information), and on the other hand the possibility to "selectively concentrate" on only a small part of the visual field. This compromise depends of course on the animal's habitat (Bertin & van de Grind 1996; see also Cliff & Bullock 1993 for a slightly different approach).

The resulting navigation is a rather uncomplicated kind of visual navigation, which we call "phototaxic foraging". Light is food and the strongest light is the most or the nearest food.

The archaepaddlers of the previous study lived in a dark environment and any light blob was regarded as food (a glowball). Thus there was no problem of segregating object and background, nor was there any need for contrast enhancement or separate ON or OFF channels[2]. In this paper we study the case of an evolving group of paddlers moving towards shallower water, where background light starts to interfere with the detection of autoluminescent prey. In the new habitat, different prey is available that reflects more light than it produces. Even nonreflecting edible objects, e.g. silhouetted against the bright surface right above our hunting paddlers, might become a meal if the paddler's visual system could develop dark-object detecting capabilities. One need not simulate the evolution process itself in this case, since it is known that the standard response of evolving visual systems to these challenges is to develop ON and OFF (and/or ON/OFF) cells and lateral inhibition to aid foreground/background segregation. One can view the Limulus visual system as a paradigm in this respect (Hartline & Ratliff 1974).

It is rather obvious, however, that this is not sufficient. One also needs to introduce circuitry to allow light and dark adaptation, so that the ON and OFF units can function satisfactorily over a wide range of background luminances, without the risk of saturation at higher luminance levels, and loss of precious information at low levels. We therefore implemented one of the simplest and most ubiquitous adaptation principles: Weber-law adaptation (see e.g. Bouman, van de Grind and Zuidema 1985, or Shapley and Enroth-Cugell 1984). An interesting property of this type of adaptation is that it emphasizes reflectancy (Shapley and Enroth-Cugell 1984). Nothing in the principles on which the resulting visual system is based is surprising or new, yet it is certainly not *a priori* clear that the new species of paddler equipped with this extremely simple diurnal visual system can successfully forage in shallow water at a variety of background light levels. In particular it is not *a priori* clear whether (and how much) lateral inhibition or light/dark adaptation contribute to foraging success. To study this question we first develop a quality-measure for the foraging behaviour (which requires side-stepping the notorious travelling salesman problem) and then study the influence of various parameters of the visual system on this quality-measure. Also the performance of the previous archaepaddler can serve as a performance-reference, to be called the "dark reference". This is based on the fact that in a simulation, one can transform the environment of a (diurnal) paddler into a deep-sea habitat by setting the background level at zero. Then the deep-sea paddler (archaepaddler or dark reference) can hunt in the transformed environment and its performance can serve as an adaptation-free and inhibition-free reference value. The importance of adaptation and inhibition can thus be appreciated by comparing the performance of a diurnal paddler with the performance of an archaepaddler in the same environment.

To keep the exposition sufficiently simple we de-emphasize the detection of objects darker than the background and concentrate on the detection of objects brighter than the background. The simulation allows us to quantify the merits for hunting success of lateral inhibition and of the automatic gain control of Weber's law. Pitting these different aspects of vision against each other is only possible in this particular kind of approach, with an explicit and biologically reasonable measure of success for the ensuing visually-guided behaviour. This opens exciting perspectives for testing theories of vision, and in another paper we will test a theory of motion vision. Here the main question is: What is the relative merit of lateral inhibition (of a simple Limulus type) and of Weber adaptation, alone or in combination, for

[2] Sets of neur(on)al connections that respond to increases or decreases in (visual) input.

simple visually-guided hunting (phototaxic navigation) in a shallow aquatic environment with variable background light levels?

2 The diurnal paddler

This new species is highly similar in many respects to its predecessor, the archaepaddler. It has a roundish body with paddles at the back and compound eyes up front. Figure 1 illustrates the network of a single cartridge[3] behind the layer of laterally inhibiting photoreceptor cells (here numbered from $i-1$ through $i+2$). Each photoreceptor acts as a Weber machine (Bouman, van de Grind and Zuidema 1985), which is described below. The i^{th} cartridge gets a central input (W_i) from the i^{th} receptor and a "surround" input (S_i) consisting of the weighted sum of W_i and the output of the nearest neighbours to the left, W_{i-1}, and right, W_{i+1}. Neurone #3 calculates the S-signal as:

$$S_i = \frac{\mathcal{C} \cdot W_i + W_{i-1} + W_{i+1}}{\mathcal{C} + 2} \tag{1}$$

This choice is convenient since it allows us to change the balance between a pure centre-drive of S (no lateral inhibition, only self-inhibition) and a pure surround-drive (no self-inhibition) by changing $\mathcal{C}$ from ∞ to 0. For $\mathcal{C} = 1$ all three inputs contribute equally to S; each one third.

FIGURE 1 ABOUT HERE.

The S-signal then passes through a leaky integrator[4] with time-constant τ_t to 'average' it a bit and thus make the filtered version of this "centre/surround" signal, $\langle S \rangle$, somewhat more sluggish than the centre signal W. Neurones #4 and #5 calculate the clipped values of $W - \langle S \rangle$ and $\langle S \rangle - W$ respectively. Clipping in this case means that the outputs (real numbers representing firing frequencies) are zero if the difference of the input signal pair is lower than a positive threshold $l \approx 0$ (for implementation-technical reasons, we actually chose $l = 10^{-5}$). If the lower clipping limit is exceeded the outputs are equal to the difference of their inputs until a saturation level of $h = 100$ is reached, the upper clipping level. This clipping applies to all neurones, and will be symbolised as clip$[x]$, with x the input to the clipping stage.

$$\begin{cases} \langle S_i \rangle = \text{RC}[S_i; \tau_t] \\ \text{ON}_i = \text{clip}\left[W_i - \langle S_i \rangle\right] \\ \text{OFF}_i = \text{clip}\left[\langle S_i \rangle - W_i\right] \end{cases} \tag{2}$$

where RC $x; \tau$ is a leaky integrator with input x, time-constant τ and a gain of 1.

The outputs of #4 and #5 thus have the character of an ON and an OFF signal, respectively. These two signals are added in neurone #6, each weighted as indicated in figure 1, to obtain the signal A which is the output-signal of the individual cartridges. Thus we combine the ON and OFF components in one ON/OFF signal for further processing. This entails no loss of information in the separate channels, while simplifying the simulations. The ON and OFF channels can always be separated if desired, as we have done in control experiments.

$$A_i = w^- \cdot \text{OFF}_i + w^+ \cdot \text{ON}_i \tag{3}$$

To ensure that the neurones in the ON/OFF channels can function over a wide range of luminances without danger for saturation, each photoreceptor-cell works as a Weber-machine measuring luminance. An analog consisting of two neurones is shown in figure 1. The simple trick is that neurone #2 divides its input signal, R, by a scaling factor that equals a constant K plus a lowpass filtered version, $\langle R \rangle$, of the input (van de Grind *et al.* 1970; Koenderink *et al.* 1970). In the Weber-adapting photoreceptor-cell, there is of course no clipping until W_i is determined!

[3] The term used to indicate the functional unit in a compound eye's retina; the circuitry behind one facet.

[4] Also called lowpass-filter; the output y to an input x is described by $y[t + \tau t] = \tau y[t] - (1 - \tau)x[t]$, with τ the time-constant.

$$\begin{cases} \langle R_i \rangle = \mathrm{RC}[R_i; \tau_w] \\ W_i = \dfrac{w_W \cdot R_i}{K + \langle R_i \rangle} \end{cases} \quad (4)$$

The lowpass filtering by a leaky integrator (#1, with time-constant τ_w) makes the feedforward gain control sufficiently sluggish, so that brief changes pass the scaler before the scaling factor can be adapted. In the special case $\tau_w = 0$, the leaky integrator looses its integrating behaviour, and the Weber-machine behaves like the well-known Michaelis-Menten saturation in enzyme-kinetics.

In its "Weber-range" (lower limit set by the constant K), slow and sustained input changes are filtered out, however, by the feedforward control that tends to keep the output of the Weber machine constant. For low background levels (below the Weber-range; i.e. $\langle R \rangle \ll K$) the feedforward path hardly adapts and thus has a constant scale factor K making the output linearly proportional to the input. A gain-factor w_W is applied to the outputs of all Weber machines.

This completes the description of the two visual modules making up a single cartridge, viz. the ON/OFF module and the Weber-machine. We will change the various parameters to emphasize the role of either of these modules. It is interesting to note here that by changing just two parameters of this light-adapting machinery, we can create a system which makes use of only the ON/OFF module. When the constant K is set to a value much higher than the average maximum light-intensity, the Weber machine acts as a linear scaler whose gain is controlled by the w_W parameter.

The output of the i^{th} cartridge is called A_i and this is the only signal used for further control of navigation. Here we use the same navigation strategy as in the archaepaddler, simple phototaxic navigation. Before describing the "command bridge" of the paddler, we need to summarise the relevant aspects of the animal's anatomy, as it evolved in the previous study (Bertin and van de Grind 1996; see also Bertin 1994).

FIGURE 2 ABOUT HERE.

Figure 2 sketches the paddler with its immobile eyes, which have an optic axis under an angle $\beta = 20°$ with the midsaggital[5] plane. The binocular overlap region is functionally important, since targets in that region are probably easier prey than those in the periphery. Thus we will give targets in that region an extra weight (ϖ) in steering control. Of the $N = 80$ cartridges that were simulated for each eye, about 54 cartridges sample the binocular region. Note that even though we model the animal in 2D, we assume that the eyes are placed on top of the animal's body. Thus occlusion by parts of the body (as suggested in figure 2) is prevented because the eyes can look over these parts. Targets entering the mouth are immediately eaten.

FIGURE 3 ABOUT HERE.

Figure 3 illustrates how visual signals from the cartridges are combined to calculate motor- commands, and we refer to this system as the (command)bridge. The simple phototaxic navigation is based on the weighted sums ($\langle A \rangle_L$ and $\langle A \rangle_R$) of the visual information (A_i) in the left and in the right eye. The weighting (the retinal weighting function, RWF) is a Gaussian around a visual axis (ε in figure 2) of optimal sensitivity, so that there is an inbuilt tendency to hunt targets in that direction unless peripheral snacks outweigh the central visual food mass. The RWF has a halfwidth σ and gain ς. Neurones in the paddler's brain clip to 0 below the lower threshold l (10^{-5}) and saturate at a level of 100.

The left and right visual signals $\langle A \rangle_L$ and $\langle A \rangle_R$ are then normalised and passed through leaky integrators (the motor neurones) which smooth fast variations and introduce a short-term memory for recent manoeuvres. This results in motor command signals which are sent to the right and left paddles, respectively. In the absence of visual information, searching behaviour is initiated by a circuit consisting of three mutually inhibiting neurones with stochastic spontaneous activity. Normally suppressed by visual information, these neurones alternatively become active, specifying bouts of swimming to the left, to the

[5] This is the plane that cuts the animal in two halves, from head to tail.

right, or straight ahead. In addition, there is a mutual inhibition between the left and right motor neurones result in a continuation and exaggeration of the last turn. Thus the paddler can continue to explore when visual information ceases, initially heading back where it just came from.

3 Methods

A. Performance indices. It has been shown that classical, positive phototaxis is an adequate strategy for foraging in a simple (deep)sea environment (Bertin and van de Grind 1996). The addition of background illumination is not expected to drastically alter this, provided an adaptation mechanism is present that functions over a wide range of luminances.

The actual foraging performance depends on the "tuning" of the animal's geometrical and physiological parameters. 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 $\wp$ (see below).

The symbols $\underline{m}$, $\underline{g}$, $\underline{s}$ and $l\underline{u}x$ will denote *arbitrary* units for length, weight, time and luminance respectively.

Three paddlers were allowed to roam a large foraging environment for a fixed period of time (4000 $\underline{s}$, with a resolution of 30 ticks per $\underline{s}$), during which data were collected. Every 1000 $\underline{s}$, eaten prey were replaced.

This procedure was repeated for each parameter value out of a relevant range. During these experiments, there was no interaction or competition between the different paddlers: thus we collect 3 independent datasets in parallel.

How do we quantify the foraging performance of a paddler? This is easily done by comparing the paddler's actual performance (the route it takes) to the optimal performance, given the paddler's starting point and the distribution of glowballs in the world. The optimal foraging route is of course found by solving the Travelling Salesman problem. Since we do not aspire to solve this problem, we define the following performance measure:

$$\wp = \frac{\mathrm{G}_e \cdot \langle \mathrm{D}_m \rangle}{\mathrm{D_t}} \cdot \frac{\langle \$_e \rangle}{\langle \$_s \rangle} \tag{5}$$

This measure basically determines the distance travelled ("cost") per eaten target ($\mathrm{G}_e/\mathrm{D_t}$). Of course, many more targets can be eaten per unit distance travelled in a high density population. Therefore, the average distance between eligible targets ($\langle \mathrm{D}_m \rangle$) is taken into account to decrease the dependency of $\wp$ on target density[6]. The measure also takes the relative caloric value of the eaten targets ("gain"; $\langle \$_e \rangle$ / $\langle \$_s \rangle$) into account. For a more detailed explanation of the meaning and computation of $\wp$ and the various symbols, we refer to the appendix.

$\wp$ approaches 1 when a paddler repeatedly takes the direct route from one target to its most profitable neighbour, where "most profitable" is the amount of food obtained per unit distance travelled. When no further information on prey distribution or routing is available, this is the optimal foraging strategy (Rössler 1974). Smaller $\wp$ values indicate lesser performance; significantly higher values are an indication that the index has become invalid (e.g. due to a too large $\langle \mathrm{D}_m \rangle$ value, or because the optimal foraging can no longer be described by $\wp$).

As argued in the Introduction, we can determine the influence of the various processes (subsystems) during the behaviour they help controlling. This can be done by defining performance measures for these subsystems as well. These can be related to the overall performance of the whole animal. We used two measures to quantify the performance of the visual system.

The first is the amount of foreground/background (FB) segregation reached by the visual system, measured as the FB-ratio $\mathcal{R}$. It is calculated as the temporal average of the ratio of the average response of all cartridges responding to foreground (prey), over the average response to background:

$$\mathcal{R} = \left\langle \frac{\sum_{i=0}^{n} A_i/n \quad \textit{; foreground}}{\sum_{i=0}^{m} A_i/m \quad \textit{; background}} \right\rangle \tag{6}$$

[6] This does not necessarily capture possible qualitative changes in optimal foraging behaviour. However, we did not observe qualitative changes over the range of target densities we used.

with $n+m=N$, and n the number cartridges whose current activity is caused by foreground (prey) plus background illumination, and m the remaining cartridges responding to background illumination only. Ideally, the latter average background response should be zero, and hence $\mathcal{R}=\infty$. In practice, transient responses can significantly decrease $\mathcal{R}$, even if after some iterations it would approach ∞.

The second subsystem performance measure quantifies the amount of relevant information present in the output of the visual system. Paddlers make use of phototaxis, i.e. they strive to balance the left and right eye-responses. To assess how much useful information is available to the paddler, we can thus calculate the temporal average of the absolute difference of the normalised eye-responses:

$$\mathcal{D}=\left\langle\left|\langle A\rangle_L-\langle A\rangle_R\right|\right\rangle \tag{7}$$

(see figure 3).

For adequate performance, it is expected that $\mathcal{D}$ should be neither too low nor too high. Too low would indicate largely identical eye-responses, due to the physical absence of visible targets, or due to an unfiltered, large background illumination. This "translates" into a paddler moving mostly straight ahead. Too high indicates the opposite; for some internal or external reason, the two eyes hardly ever agree. The result is of course, a paddler continuously making large course alterations.

Obviously, $\wp$, $\mathcal{R}$ and $\mathcal{D}$ are not available to the simulated paddlers; only to us as external observers, looking over the paddler's shoulder, and evaluating each and every of its decisions and moves.

B. Simulation Methods. Experiments were carried out using a proprietary simulation package written in ANSI C, and run on HP 9000/730, Silicon Graphics 4D and Apollo DN10000 computers. We used a time-step of $\delta t=1/30s$; on the HP, a simulation second takes 0.022 to 0.16 real seconds per paddler. In all experiments, paddlers were "released" in a foraging space (*world*) of 150 $\underline{m}$ square, within which $\mathbb{G}=121$ targets were uniformly distributed. The prey had a radius of 0.5 ± 0.144 $\underline{m}$ (taken from a uniform distribution). Prey autoluminance was fixed at $5\times10^{-4}\pm2.88\times10^{-5}$ $l\underline{u}x$, with a reflectance of $0.05\pm2.88\times10^{-3}$ (both uniform). Thus targets are visible in both nocturnal and diurnal circumstances; the effective luminance of a glowball i is its autoluminance ($L_g[i]$), plus its reflectance ($\mathcal{R}[i]$) times the background illumination ($L_b[x,y]$) at its location:

$$L[i]=L_g[i]+\mathcal{R}[i]\cdot L_b[x,y] \tag{8}$$

The paddlers had a radius of 2 $\underline{m}$ and a mass of 500 $\underline{g}$. For a full listing of parameter-values giving "standard" performance, we refer to Bertin and van de Grind (1996). Parameters relevant to the current discussion are listed in the captions to figures 2 and 3. The threshold for switching to searching behaviour, and the scotopic threshold[7] were set to 10^{-5}.

Paddlers venturing too far from (i.e. loosing visual contact with) the world were replaced at a random position within the world. This is merely a trick to cope with the problems posed by using a finite world (instead of using a toroidal world, which would be more complicated in terms of implementation). The transient effects of "teleporting" a paddler to a new position are negligible for the simulation time used. It also gives a penalty to paddlers that dwell too long in the boundary zone where they have not lost visual contact completely. Note that background illumination does extend beyond the boundaries of the world (which would not be possible in a toroidal world).

Background illumination is modelled by a function which defines the local amount of background illumination at a given location in the world. This amount is added to the illumination an object (e.g. an eye) receives at that location.

Two different types of background illumination were used. In case of random (*rd*) illumination, each 1×1 $\underline{m}$ square has a luminance value taken from a random distribution between 0 and $\mathcal{L}_{\max}$ $l\underline{u}x$. These random luminances vary once every simulated second; the regime is used as a most demanding regime. A ramp (*rp*) illumination indicates a continuous luminance gradient from 0 $l\underline{u}x$ in the upper-right corner, to $\mathcal{L}_{\max}$ $l\underline{u}x$ in the lower-left corner of the world. This regime emulates the effect of a sloping sea-floor, or light blocked by overhanging vegetation. In both cases, a modulation (of the form $\cos[1.59\times10^{-3}\cdot(x-y)]$, with a modulation depth of 6.75% of the average background luminance) is added to simulate e.g. the

[7] The lowest luminance level where the visual system still works

effect of waves. Thus, in the case of a ramp illumination, the background illumination at (x, y) is given by:

$$L_b[x, y] = \text{ramp}[x, y] + 0.0675 \cdot \frac{\mathcal{L}_{\max}}{2} \cdot \cos\left[1.59_{\times 10^{-3}} \cdot (x - y)\right] \quad (9)$$

where ramp defines a gradient which is 0 at $(0, 0)$ and $\mathcal{L}_{\max}$ at (x_{max}, y_{max}).

When not filtered out, a ramp illumination will cause a paddler to swim either up (stable solution) or down (unstable) the gradient, as it strives to equal the response of the two eyes, turning towards the eye with the highest response. In an environment with sufficient random background illumination, any paddler will perform a random walk dictated by the background light distribution.

4 Experiments and results

A. Internal properties: time-constants and lateral inhibition What is the influence of the leaky-integrator time-constants in the ON/OFF channels and the Weber machine? Let us first discuss the influence on the visual performance measures. Larger time-constants will of course result in longer delay and more memory effects in these integrators. This in turn means that the ON/OFF channels are slower in following their input, and that the Weber-machine is slower in adapting to changing levels of ambient light. In other words: the effect of inputs lasts longer.

We should study the time-constant of the ON/OFF channels independent of the amount of lateral inhibition ($\mathcal{C}$), given the fact that this parameter determines the inputs to the ON/OFF channels. Therefore, we also varied $\mathcal{C}$ over its complete range from 0 to ∞ — pure lateral inhibition to no lateral inhibition.

Figure 4 shows how variation of these parameters influences the overall performance of the diurnal paddler.

FIGURE 4 ABOUT HERE.

Paddlers were tested in a ramp background illumination regime with $\mathcal{L}_{\max} = 150 lux$. This is an intermediate value where the Weber-machine is in its adaptation range at the bright end of the ramp. On average, the paddlers experienced a background illumination of $86 \pm 25 lux$. The time-constant τ_w of the Weber machine was varied between $\tau_w = 0$ and $\tau_w = 2.3$. The time-constant τ_t of the ON/OFF channels was covaried with τ_w. (see the caption to figures 4a and 4b for more details)

At the low end of the τ_w scale, the average background response is zero, so $\mathcal{R} = \infty$. For higher τ_w and τ_t, the foreground response increases due to memory effects. The background response increases faster so $\mathcal{R}$ drops. Decreasing the strength of lateral inhibition (i.e. higher $\mathcal{C}$ values) reduces the contrast-enhancement, and increases the sensitivity to global luminance fluctuations. Both effects result in a lower $\mathcal{R}$. In all cases, the decrease of $\mathcal{R}$ levels off above a certain τ_w respectively $\mathcal{C}$ value. For the lowest τ_w, adaptation in the ON/OFF channels can be so fast that the transient responses to temporal variations are too short to measure. If in that case there is no lateral inhibition ($\mathcal{C} = \infty$), the paddler is effectively blind, and $\mathcal{R}$ is not defined (the dotted trace for $\mathcal{C} = \infty$ in figure 4a starts only at $\tau_w = 0.015$). Here, the average difference between the normalised eye-responses ($\mathcal{D}$) is zero (not shown). In the other cases, $\mathcal{D}$ increases with increasing τ_w; longer delays in adaptation (more memory effects) cause more and more unequal responses in the two eyes. $\mathcal{D}$ also increases with $\mathcal{C}$, which causes a larger sensitivity to variations in background illumination. When the time-constants become too long, a target projection moving over the retina leaves an increasingly intense set of after-images, which are counted as background response; $\mathcal{R}$ decreases). An "after-image" in the ON/OFF channel is in fact the 2nd derivative of the luminance input, the Weber-response taking a 1st derivative. Therefore, an ON luminance step results in a ON response, followed by a (longer) OFF response. Judging from direct observations of the retinal activities, the effect of after-images is especially strong for target-projections moving at high speed across the retina. Even contrast-enhancement by lateral inhibition is no longer of any help: the cross-connections cause a lateral spreading of the after-image to the two neighbour cartridges.

How do these changes in vision work out on the foraging performance as indicated by $\wp$? We find a decreasing performance with increasing τ_w and $\mathcal{C}$ that eventually levels off. An exception are of course the blind paddlers without any lateral inhibition ($\mathcal{C} = \infty$) and very short time-constants: these have a

very low performance, as is evident from figure 4b. It can also be seen that increasing the time-constants in these paddlers ameliorates their performance somewhat ($\mathcal{R}$ increases; cf. figure 4a). due to increasingly intense after-images, performance stops increasing (and $\mathcal{R}$ decreases) when the time-constants become too long. It is to be expected that for time-constants still larger than our upper value, performance will also degrade.

We also performed the experiment for relatively longer ON/OFF time-constants ($\tau_t = \frac{10}{3}\tau_w$), and for a fixed $\tau_t = 0.05$[8]. In both cases, the influence of the time-constant(s) and/or the amount of lateral inhibition is less apparent. In the former case, the shortest time-constants give a much higher $\mathcal{R}$, but without a better foraging performance. In the latter case, performance $\wp$ is almost constantly good. This is an indication that over the range of time-constants we studied, the time-constant τ_t of the ON/OFF channel is the one with the most influence, and that the effect of lateral inhibition is most prominent at lower τ_t.

B. External properties: background luminance level The results presented above show that the diurnal paddler can perform adequately in an environment with an intermediate level of ambient light. In this section, we will examine how foraging performance varies with varying amounts of ambient light.

Based on the results presented above, we chose a parameter setting $\mathcal{C} = 1$ and $\tau_w = \tau_t = 0.05$. The resulting paddler was subjected to a large range of luminance levels in order to test its adaptive potential. In murky coastal waters, a vertically migrating animal has to rely on light generated by bioluminescence, even at relatively small distances below the surface: absorption can be so strong that daylight does not reach deeper than 60m (Lythgoe, 1979). At the other end of the scale, in extremely clear water, a substantial part of the spectral frequencies present in daylight is hardly absorbed at all in the upper few meters.

We thus studied performance without background illumination ($\mathcal{L}_{\max} = 0$), and in a ramp illumination regime with an intensity ranging from $\mathcal{L}_{\max} = 0.05$ to $\mathcal{L}_{\max} = 50000 l\underline{u}x$. As a reference, an archaepaddler was subjected to the same range of luminances: the performance of this paddler in complete darkness (with the targets visible through their autoluminescence) is referred to as the dark reference. Figures 5 and 6 show the $\mathcal{D}$ and $\wp$ measured as a function of the background illumination averaged over the paddlers' paths.

FIGURE 5 ABOUT HERE.

The average difference between the normalised eye-responses ($\mathcal{D}$: see figure 5) remains fairly constant at dark reference level. When the average background illumination approaches the constant K of the Weber machine (compare trace 2 with traces 1 and 3), Weber adaptation begins to take place, and as a result, the difference in response between the eyes increases. At a luminance-level approximately 100 times higher, $\mathcal{D}$ becomes constant. Its high value at these luminance levels is caused mostly by the fact that the more variable transient response to a moving target-projection is much larger than the tonic[9] responses to the edges of that target-projection. Also, a very small response to background illumination is noticeable. As the two eyes will generally not receive an identical amount of background illumination, this tends to increase $\mathcal{D}$, especially when no targets are in sight. In a random illumination, $\mathcal{D}$ increases much faster.

In the reference paddler, $\mathcal{D}$ shows a different behaviour that is purely dictated by the background illumination. Thus for non-zero $\mathcal{L}_{\max}$, $\mathcal{D}$ also increases in a random illumination (thick trace, open diamonds in figure 5), but it decreases in a ramp illumination (thick trace, filled circles). When the average illumination approaches $8 l\underline{u}x$, $\mathcal{D}$ drops in both illumination regimes. This is a result of neural saturation (80 photoreceptors, with a gain of 10, and a maximal frequency of 100) in the "normalising centre" of figure 3.

The importance of the Weber machine in preventing neural saturation can be appreciated by comparing the results for $w_W = 1$ and $w_W = K$ (curves 1 and 3 in figures 5 and 6). $w_W = K$ means that

[8] the time-constant reported for the highpass filter which resides in the fly LMC — Large Monopolar Cell, a rather front-end "processing unit" which seems to be involved in dark/light adaptation (Egelhaaf & Borst 1992).

[9] tonic = lasting ; phasic = transient, passing

the theoretical fully adapted state of the Weber machine equals the maximal firing frequency; a sufficient increase in illumination would then cause saturation. Yet this does not happen to a significant degree.

The separation index $\mathcal{R}$ (which decreases for $\mathcal{L}_{\max} > 0$) also levels off when $\mathcal{D}$ becomes constant. Its final value is still between 3 and 5 times higher than that of the reference paddler.

FIGURE 6 ABOUT HERE.

Finally, let us look at the foraging performance: see figure 6. The diurnal paddler performs at reference level for all but the higher background luminances, where a slight drop can be observed. At those luminances, the background increasingly controls the paddler's manoeuvres. Note that the much earlier onset of Weber adaptation for $K = 1, w_W = 1$ does not result in significantly worse performance. On the other hand, performance is worse ($\wp$ is significantly lower than its "ideal" value 1) at low luminances for $K = 100$ and $\varsigma = 1$. Here information is lost due to an insufficient gain in the Weber machine.

C. Decrement detection: effect of the OFF channels How is the above picture altered by assigning a non-zero weight to the OFF channels? To answer this question, we repeated the experiments with $w^- = w^+ = 1$. This means that both ON and OFF channels have equal votes in the paddler's visually-guided behaviour.

The most significant effect is that the responses to both foreground and background illumination are more or less doubled. In some cases, like $K = w_W = 1$, this causes a small increase in the foreground/background separation ($\mathcal{R}$) for the higher background luminance levels. However, $\mathcal{R}$ decreases for zero background illumination: a result of after-images in the OFF channels. These effects do not result in significantly different overall performance ($\wp$).

A clearer effect can be found for other combinations of the time-constants and the amount of lateral inhibition. For strong lateral inhibition (low $\mathcal{C}$) and small time-constants, $\mathcal{R}$ increases with respect to the case $w^- = 0$, while it decreases for lower lateral inhibition and/or higher time-constants. These effects are also visible in the overal performance. They can be explained by a stronger (respectively weaker) tonic response to the edges of target-projections, which provide a more solid foundation for visual decisions amidst the much more variable transient responses.

5 Discussion

We have studied simple visually-guided foraging in a diurnal environment. Ambient light provides a means for detecting targets (from reflection), but it also introduces the need to "discount the background" in phototaxic foraging. A mechanism is needed to distinguish object illumination from background illumination and to do this reliably over a wide range of background intensities. We therefore borrowed the ideas of Weber adaptation and lateral inhibition from the literature on vision research and studied the question whether one of them or both together suffice(s) to solve the problem of object segregation under a wide range of ambient light intensities. At a given level of ambient light, we find a rather large range of time- constants and lateral inhibition ratios for which adequate ("near-optimal") behaviour results, both in terms of foreground/background segregation at the retinal level, and in terms of the paddler's foraging performance. In general, we find that performance is best for short time-constants, and a small lateral inhibition fraction. The latter means that the ON/OFF channel compares the centre of its receptive field to a measure of the activity in its entire receptive field, where the surround weighs at least as much as the centre. As a result, the paddler effectively detects the edges of targets only. Edges ("spatial variations") give rise to a tonic response. Low time-constant values mean that both filters in the adapting system respond to changes with minimal delay, and do not "smear out" events in time. Temporal variations thus give rise to a short transient response.

With weak lateral inhibition (large lateral inhibition ratios), only short to very short time-constants give rise to adequate performance. With no lateral inhibition at all, more temporal integration is needed, and mediocre performance results at best for long time-constants. The importance of lateral inhibition can best be seen at intermediate levels of background illumination, where the Weber machine's transition from a passive scaler to an adaptive scaler takes place. In this range, lateral inhibition is crucial to filter out

global luminance information that would otherwise get through. For stronger background illumination, the Weber machine's response becomes transient, responding only to the changes in retinal illumination caused by target projections moving across the retina. In short, dark/light adaptation in this model is best when local changes are detected with respect to (more) global constancy. In this way, the presence of a target can be detected against the background. Using this well-known principle, we find that over a substantial range of ambient light levels, foraging performance can be on a par with the performance of an archaepaddler in its own habitat (the dark reference, defined in the introduction). The limits of this range of good performance depend on the scaling constant and on the gain of the Weber machine. The scaling constant defines a level of input above which the Weber machine makes a transition from a passive filter (scaling with a constant) to an active, adaptively scaling filter.

A high ambient light level can cause large responses in the ON and/or OFF channels. A large scaling constant will then result in better performance. However, depending on the gain, the Weber machine's output might in that case become sub-threshold for low levels of ambient light, or complete darkness, resulting in sub-optimal performance. Therefore a large scaling constant with a low gain may be beneficial for life at very high luminance levels, but becomes a handicap at lower light levels. On the other hand, too low a scaling constant is suboptimal since adaptive scaling can then be initiated too soon. This means that no efficient use is made of the nervous system's bandwidth, and that precious information can get lost. A compromise must (and can) be found between these two extremes if hunting is necessary both at high and at low background levels. With the given model, detection of prey is possible over a wide range of ambient light levels, using just one type of photoreceptors. This range can be tuned to the paddler's need by adjusting just 3 parameters of the light adapting system (where the two parameters of the Weber machine can have identical values). In this way, paddlers that live in absolutely dark to (moderately) light environments can make do with the same mechanism that allows survival in moderately light to very bright environments.

An interesting finding is that a dark/light adapting system with the best foreground / background segregation does not necessarily warrant the best foraging performance. Optimal foreground/background separation for intermediate levels of ambient light occurs e.g. for small time-constants, without lateral inhibition. Yet in that case performance is very bad, being solely based on blind search: all visual information is filtered out so fast that visual guidance of the behaviour is impossible. On the other hand, good performance can result within a range of foreground/background response ratios ranging from as low as approximately 5 up to values several orders of magnitude higher. Similarly, the OFF channel does not contribute to the information carried by the ON channel in such a way that it significantly improves overall performance. For some settings, performance is even worse! Therefore, given the task of detecting objects reflecting ambient light, it seems better not to use OFF channels. This changes, of course, if dark objects (intercepting ambient light) become important either as prey or as predator. It therefore appears that OFF channels find their use exclusively in that realm.

A side-effect of dark/light adaptation by the present mechanism is that prey are detected by their edges only. This means that size information, indicating distance to, or "electability" of the prey, is no longer present in terms of the total light response. With the current simple foraging strategy, this is not a problem, but it can become problematic when more elaborate foraging strategies arise. In that case it is necessary to process contrast and edges rather than merely "light" information. In principle, size information can be retrieved by a matching of the edges and comparing their local sign value, or by propagating neural activity ("filling in") inside matched edges. In the present 2D implementation of the model, i.e. with a 1D retina, correct matching of the edges is problematic. In a 2D retina belonging to a 3D model, the edges belonging to one target can in principle form a closed (elliptic) pattern of activity and thus be lumped. This would lead to a kind of silhouette vision. Rather then going in this direction of more sophisticated object analysis it would seem to be more profitable for foraging success to better analyse where objects are and where they are going. Anything of the proper size that can be hunted down can then be tasted or courted. The same ON/OFF channels of the above system can be connected in a different way to form retinal image motion detectors similar to those reported in the fly (Franceschini 1989, Egelhaaf and Borst 1992, 1993). Such elementary motion detectors are a necessary condition for extracting information from the optic flow, such as information on speed and direction of egomotion, and time to contact to external objects. Paddlers with such enhanced capabilities of visual analysis will be described elsewhere.

From the above findings it is clear that the visual sytem should be tuned to the task and environment in order to reach satisfactory performance. This emphasises the need for an approach to vision as described in this paper, where a complete perception-action cycle is studied rather than the visual system in isolation. The parameters of the visual system might depend more strongly on the requirements of interacting with the world than on requirements to optimally "see" the world.

Given this insight, it is interesting to note that another way to study such dependencies is by use of (unsupervised) evolutionary methods, e.g. genetic algorithms. Presently, efficient techniques for "coding" neural networks in genomic information are becoming readily available (e.g. Gruau 1994). This makes it feasible to embark on this kind of studies (see e.g. Beer & Gallagher 1992 or Cliff *et al.* 1993), which can provide unexpected new results. It should be noted, however, that the parametric studies presented in this paper will continue to be useful to analyse the results (e.g. those in Cliff *et al.* 1993) of such unsupervised evolution. Just as the result of natural evolution...

6 Appendix 1: a foraging performance measure, $\wp$

One performance measure might be the path choice efficiency. If the minimal total distance that has to be travelled to optimally fill the stomach is $\mathrm{D}_{\mathrm{min}}$, and the actual total distance travelled to reach this goal is D_{t}, then the ratio $\wp_p = \mathrm{D}_{\mathrm{min}}/\mathrm{D}_{\mathrm{t}}$ is obviously a measure of the path choice efficiency, where optimal performance leads to $\wp_p \mapsto 1$.

Now a paddler that only optimises path choice efficiency might not serve its goal of quickly filling its stomach optimally. If it could reach the same end state with fewer prey of larger size it might be better off. However, the influence of target size does not show up in $\wp_p$. This is so, since both $\mathrm{D}_{\mathrm{min}}$ and D_{t} increase with the number of eaten targets, irrespective of their size. The ratio is about constant for increasing numbers of (smaller) prey that are visited with the same navigation efficiency. Thus we need an additional factor for target choice efficiency, $\wp_t$, which can be defined in an obvious way as the ratio of the average caloric value ('size') of the eaten prey, $\langle \$_e \rangle$, over the average caloric value ('size') of all prey present at the beginning of foraging, $\langle \$_s \rangle$. Thus $\wp_t = \langle \$_e \rangle \,/\, \langle \$_s \rangle$. Again an optimal target choice would tend to lead to $\wp_t \mapsto 1$. An overall performance factor $\wp$ can be defined as the product of $\wp_p$ and $\wp_t$, and thus takes both path choice and prey choice into account.

Of all mentioned terms, $\mathrm{D}_{\mathrm{min}}$ is the most problematic, since its calculation requires that we solve the Travelling Salesman Problem, a feat of which we are not capable without putting constraints on the glowball population. To approximate $\mathrm{D}_{\mathrm{min}}$ for any given glowball population, we therefore resort to a simpler strategy. First we estimate the average $\langle \mathrm{D}_m \rangle$ of all distances to the most profitable next target from the position of each prey ("eating positions"), in a way described below in algorithmic fashion. Then we approximate $\mathrm{D}_{\mathrm{min}}$ as the product of the number of eaten prey during the whole foraging run and this (estimated) average minimal distance: $\mathrm{D}_{\mathrm{min}} = \mathrm{G}_e \cdot \langle \mathrm{D}_m \rangle$. Taken together, this means that we define the performance index as shown in formula 5 of the main text.

To account for changes in the prey distribution during foraging (e.g. as a result of eating), $\langle \mathrm{D}_m \rangle$ and $\langle \$_s \rangle$ are updated at regular intervals during a simulation.

The average minimal distance $\langle \mathrm{D}_m \rangle$ in a prey distribution is determined by an algorithm that takes information from both the prey distribution (caloric value & distance) and the paddler's capabilities (its field of view) into account. To account for variations in the prey distribution, $\langle \mathrm{D}_m \rangle$ is calculated several times during each simulation run. In the following pseudo-code description, n stands for the number of times $\langle \mathrm{D}_m \rangle$ has been calculated in the current run.

1. if `n == 0`, then `total_distance = 0` and `segments = 0`.
2. take a target, `A`, from the population, G.
3. if `A` is known to be the "`nearest_of`" to another target, `C`, then the orientation of an observer `P` moving from `C` to `A` is defined, and thus `P`'s field of view, φ, can be taken into account. (This information is not available for all targets during the first run on a new population (`n == 0`), and may be outdated when the target distribution changed.)

4. find a target, `B`, among all other targets (including `C`) which compromises between a minimal cost (`distance(A to B)`) and a maximal `reward` (caloric value; size times luminance of the target): minimise `distance/reward`. If φ is defined, also find the target `B*` with minimal `distance/reward` and visible by `P` (within φ).

5. if a target `B*` has been found, then `B = B*`.

6. tell `A` that, from its position, the most profitable next target is `B`:
   `A->nearest_to = B`.
   Also tell `B` that it is `A`'s "`nearest_to`" target:
   `B->nearest_of = A`.
   Update `total_distance` and the number of segments:
   `total_distance += distance(A to B) ; segments += 1`.

7. repeat from 2 for all targets in $\mathbb{G}$.

8. update $\langle \mathbb{D}_m \rangle$ and the counter `n`:
   $\langle \mathbb{D}_m \rangle = \frac{\texttt{total\_distance}}{\texttt{segments}}$ `; n += 1`.

Acknowledgements

The research presented in this paper was part of a PhD project of the first author 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).

We thank dr. A. Noest and dr. M. Lankheet for critically reading the manuscript, and the BioInformatics department (Prof. P. Hogeweg) for the use of their workstations.

References

Beer R.D. (1990) *Intelligence as Adaptive Behavior - An Experiment in Computational Neuroethology*. Boston, Academic Press 1990

Beer R.D. and Gallagher J.C. (1992) *Evolving Dynamical Neural Networks for Adaptive Behavior*. Adaptive Behavior 1 1992: 91-122

Bertin R.J.V. (1994) *Natural smartness in Hypothetical animals-Of paddlers and glowballs*, PhD thesis

Bertin R.J.V. and van de Grind W.A. (1996) *Phototropic foraging in the archaepaddler, a hypothetical deep-sea species*. Submitted for publication.

Bouman M.A., van de Grind W.A. and Zuidema P. (1985) *Quantum fluctuations in vision*. In: *Wolf W. (ed): Progress in Optics Vol.* XXII. North-Holland, Amsterdam 1985: 39-144

Braitenberg V. (1984) *Vehicles - Experiments in Synthetic Psychology*. Cambridge, MA, The MIT Press 1984

Cliff D. and Bullock S. (1993) *Adding "Foveal Vision" to Wilson's Animat*. Adaptive Behavior 2 1993: 49-72

Cliff D., Husbands P. and Harvey I. (1993) *Explorations in Evolutionary Robotics*. Adaptive Behavior 2 1993: 73-110

Corbacho F.J. and Arbib M.A. (1995) *Learning to Detour*. Adaptive Behavior 4 1995: 419-468

Cruse H., Brunn D.E., Bartling Ch., Dean J., Dreifert M., Kindermann J. and Schmitz J. (1995) *Walking: A Complex Behavior Controlled by Simple Networks*. Adaptive Behavior 4 1995: 385-418

Dennett D.C. (1995) *Darwin's dangerous idea*. Simon & Schuster, New York 1995

Egelhaaf M. and Borst A. (1992) *Are there seperate ON and OFF channels in fly motion vision?*. Vis. NeuronSci. 8 1992: 151-164

Egelhaaf M. and Borst A. (1993) *Motion computation and visual orientation in flies*. Comp. Biochem. Physiol. 104A 1993: 659-673

Ekeberg O., Lansner A. and Grillner S. (1995) *The Neural Control of Fish Swimming Studied Through Numerical Simulations*. Adaptive Behavior 4 1995: 363-384

Franceschini N., Riehle A. and le Nestour A. (1989) *Directionally Selective Motion Detection by Insect Neurons*. In: *Stavenga & Hardie (eds.): Facets of Vision*. Berlin/Heidelberg, Springer-Verlag 1989: Ch. 17, 360-390

Gibson J.J. (1979) *The ecological approach to visual perception*. Houghton Mifflin Company, Houston

van de Grind W.A., Koenderink J.J. and Bouman M.A. (1970) *Models of the Processing of Quantum Signals by the Human Peripheral Retina*. Kybernetik 6 1970: 213-227

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

Gruau F. (1994) *Automatic Definition of Modular Neural Networks*. Adaptive Behavior 2 1994: 151-183

Hartline H.K. and Ratliff F. (1974) *Studies on excitation and inhibition in the retina*. Chapman and Hall, London

Koenderink J.J., van de Grind W.A. and Bouman M.A. (1970) *Models of Retinal Signal Processing at High Luminances*. Kybernetik 6 1970, 227-237

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

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

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

Rössler O.E. (1974) *Adequate locomotion strategies for an abstract organism in an abstract environment-A relational approach to brain function*. In: *Conrad M., Güttinger W. and Dal Cin M. (1974): Lecture Notes in Biomathematics 4: Physics and Mathematics of the Nervous System*. Berlin/New York, Springer Verlag 1974: 342-370

Simmons P.J. (1993) *Adaptation and responses to changes in illumination by second- third-order neurones of locust ocelli.* J. Comp. Physiol. 173: 635-648

Shapley R.M. and Enroth-Cugell C. (1984) *Visual adaptation and retinal gain controls*. Progress in retinal research 3 1984: 263-346

Terzopoulos D., Tu X., and Grzeszczuk, R. (1994) *"Artificial Fishes with Autonomous Locomotion, Perception, Behavior, and Learning in a Simulated Physical World.* In: *Brooks R. and Maes P. (eds.): Artificial Life IV: Proc. of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, MA, 1994: p.17-27

Walter W.G. (1950) *An imitation of life*. Scien. Am. 182: 42-45

Walter W.G. (1951) *A machine that learns*. Scien. Am. 185: 60-63

7 Captions and Figures

Figure 1 :

The diurnal paddler's retina. A lateral inhibition layer feeds $N = 80$ cartridges, each containing a Weber machine working on a weighted sum of ON and OFF channels. $\rightarrow$: a "normal" excitatory connection; —┤: an inhibition; —○: a shunting inhibition (here with the scaling constant K). ⓿ is used to indicate a dividing neurone; Ⓢ indicates a summator; Ⓧ a multiplier and Ⓣ a leaky integrator. The photoreceptors have a gain of 10.0. The white box shows the clipping function performed by all cells: $l = 10^{-5}$ and $h = 100$.

Figure 2 :

The paddler's anatomy. The two eyes were placed at a position of $\alpha = 90°$, sampling an angle of $\varphi_v = 140°$, centered around the eye-axis at $\beta = 20°$. Light absorption by the water was approximately 0.125 $l\underline{u}x/\underline{m}$ and the direction of optimal sensitivity was at an angle $\varepsilon \approx -23.18°$ to the eye-axis. Thus the effective visual horizon was at $\mathbb{D}_v = 26\ \underline{m}$, and $\mathbb{D}_\varepsilon = 36\ \underline{m}$. The inset shows a screendump of an actual paddler, in which all proportions are correct. Negative angles indicate clockwise rotations. Values given are for the left eye: for the right eye, multiply angles by -1.

Figure 3 :

Visuomotor network of the diurnal paddler's nervous system.The weighted sums of the A_i responses (2 times N) from the left and right eye are projected onto respectively the right and the left paddles, after having been normalised and temporally averaged by leaky integrators. In the absence of visual information, searching behaviour is generated (circuit not shown). Weighting of the A_i responses is Gaussian, with a halfwidth $\sigma = 35°$ centered around $\varepsilon = -23.18°$ (see figure 2), and height w_W. Individual A_i responses from the binocular field receive an additional weight-factor $\varpi = 0.5$. Clipping parameters as in figure 1.

Figure 4a :

$\mathcal{R}$ as a function of $\tau_w$and $\mathcal{C}$.The average foreground/background ratio as a function of the lateral connection weights and the (coupled) time-constants with $K = 100$, $w_W = 100$ and $\tau_t = \tau_w$, for a ramp illumination regime with a luminance range of 150. Performance is ideal for $\tau_w \leq 0.001$. Performance progressively decreases with increasing $\mathcal{C}$ and increasing τ_w; in both cases, the decrease of $\mathcal{R}$ saturates. Numbers in the curves indicate the $\mathcal{C}$ values; error-bars indicate standard deviations..

Figure 4b :

$\wp$ against $\tau_w$and $\mathcal{C}$.Performance as a function of the lateral connection weights and the (coupled) time-constants with $K = 100$ and $w_W = 100$, for a ramp illumination regime with a luminance range of 150. Performance is best for low τ_w and low $\mathcal{C}$. Note that for $\mathcal{C} = \infty$ (only the central photoreceptor's output is used), the paddlers perform in a mediocre way at best, for higher τ_w only. In this case, the tonic response due to spatial filtering is absent. Therefore almost all visual information is discarded (the paddlers are "in search mode" almost constantly), except for the larger τ_w values which introduce their own disadvantages. Numbers in the curve indicate the $\mathcal{C}$ values; error-bars indicate standard deviations.

Figure 5 :

$\mathcal{D}$ against average background illumination,measured during a simulation similar to the one in figure 4. All diurnal paddlers shown here remain close to the reference performance of an archaepaddler at $\mathcal{L}_{\mathsf{max}} = 0$ for at least a fraction of the range of illuminations that were experienced. For average background illuminations approaching K, Weber adaptation occurs. This is visible as an increase in $\mathcal{D}$. Neurones saturate at a firing level of 100 (spikes per timestep). $\tau_w = \tau_t = 0.05$, $\mathcal{C} = 1$. Trace 1: $K = 100, w_W = 1$; trace 2: $K = w_W = 1$; trace 3: $K = w_W = 100$; trace 4: reference paddler, ramp illumination; trace 5: reference paddler, random illumination. Error-bars indicate standard deviation.

Figure 6 :

$\wp$ against average background illumination;performance measured during the simulation shown in figure 5. It is clear to see that performance of all diurnal paddlers remains on the dark reference level for the complete range of background illuminations, despite the increase in $\mathcal{D}$ visible in figure 5. Note that, at low background illumination,

when it can still more or less separate targets from the background, the archaepaddler performs better in a ramp regime. However, at high background illumination, performance is somewhat better in a random regime. Now the distraction of the randomly changing background completely guides the paddler, preventing it from swimming in straight paths. Error-bars indicate average of the standard deviation of each curve. Trace 1: $K = 100, w_W = 1$; trace 2: $K = w_W = 1$; trace 3: $K = w_W = 100$; trace 4: reference paddler, ramp illumination; trace 5: reference paddler, random illumination. Error-bars indicate standard deviation.

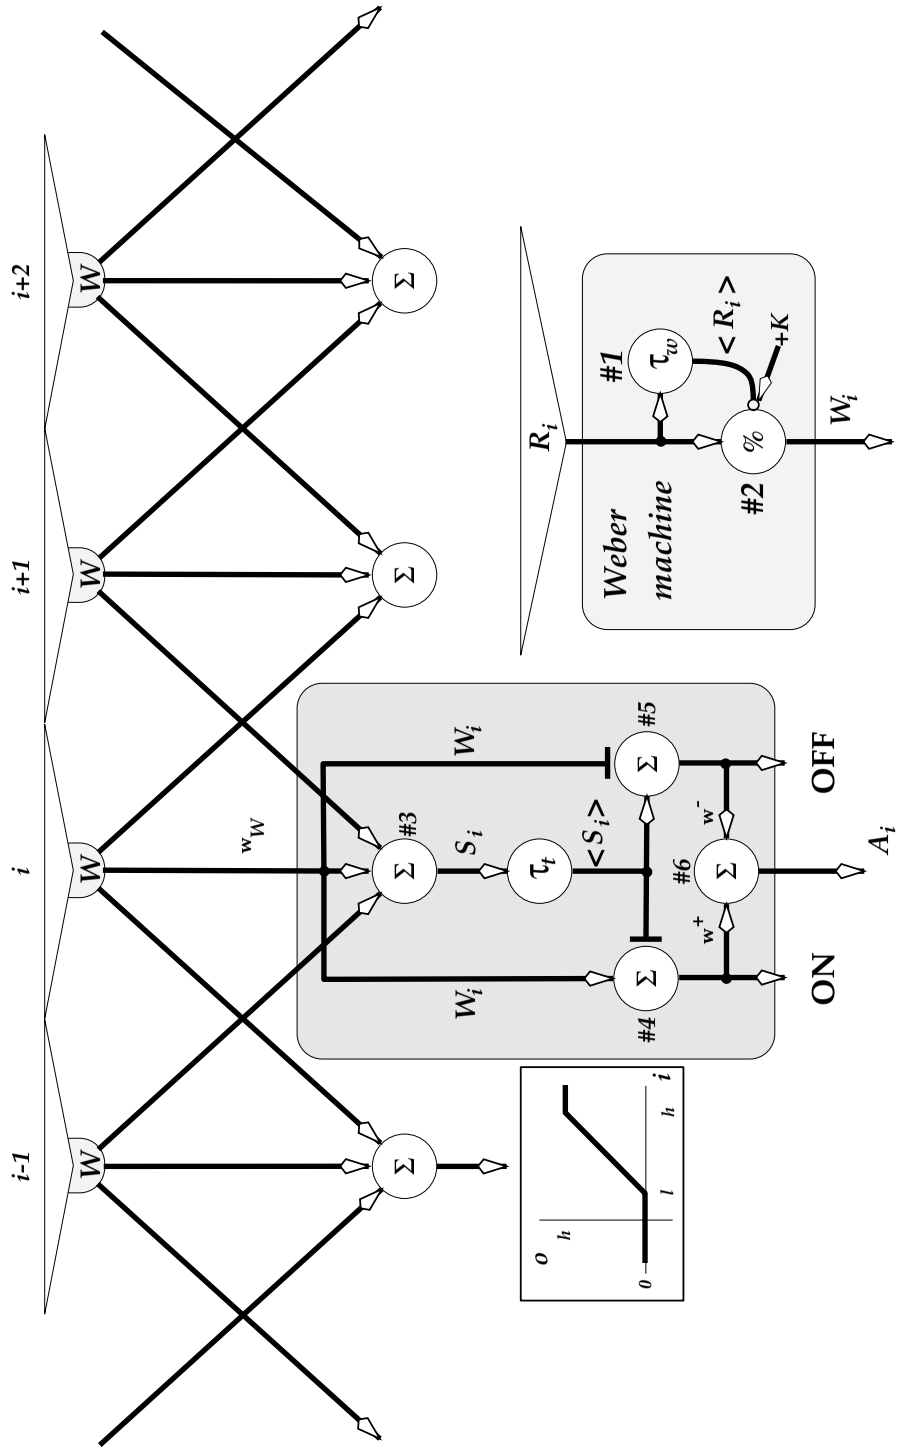
i-1
i
i+1
i+2
W
Σ
R_i
Weber machine
#1
τ_w
$< R_i >$
+K
#2
%
W_i
w_W
#3
S_i
τ_t
$< S_i >$
#4
#5
#6
w^+
w^-
ON
OFF
A_i
o
h
l
i

Figure 1:

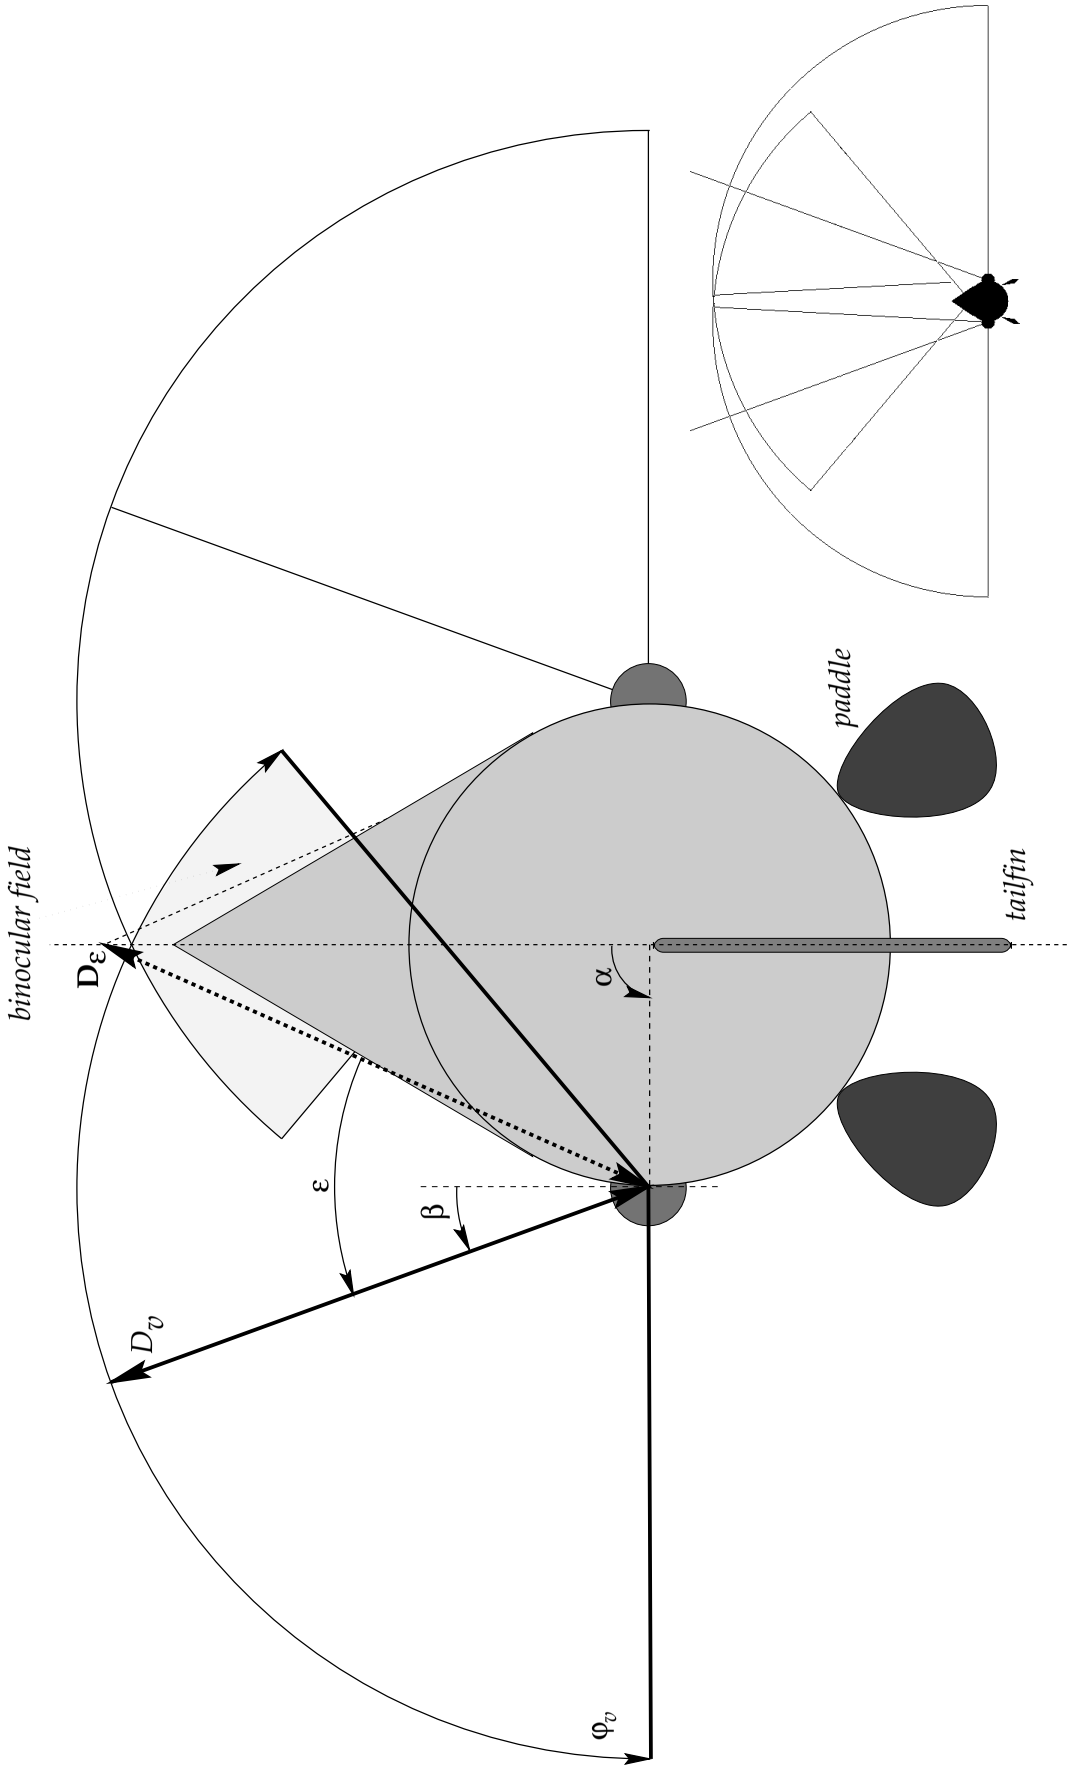

Figure 2:

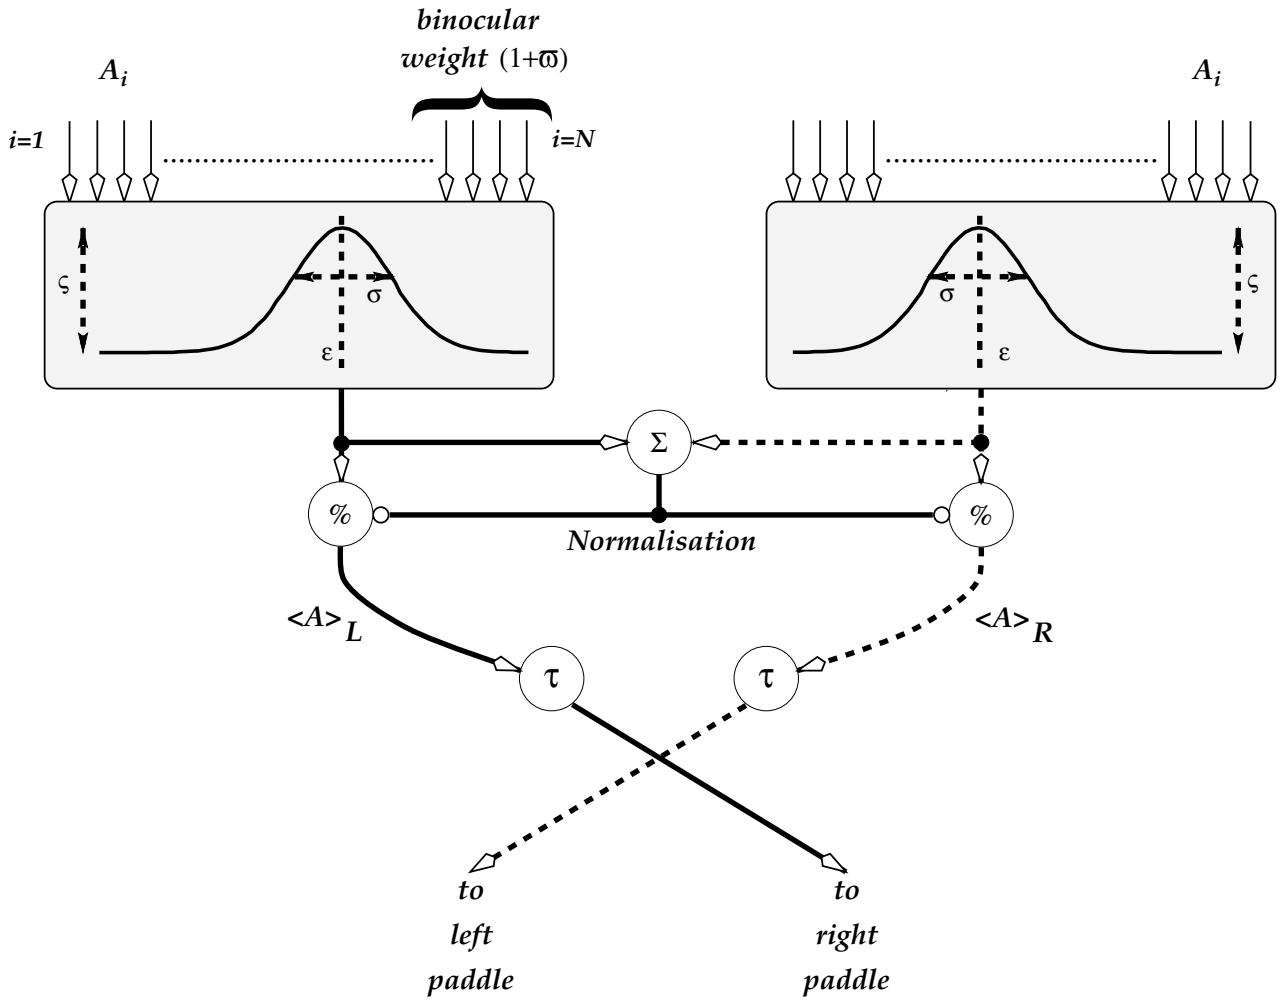
binocular
weight (1+ϖ)
A_i
A_i
i=1
i=N
ς
σ
ε
σ
ε
ς
Σ
%
%
Normalisation
<A>L
<A>R
τ
τ
to
left
paddle
to
right
paddle

Figure 3:

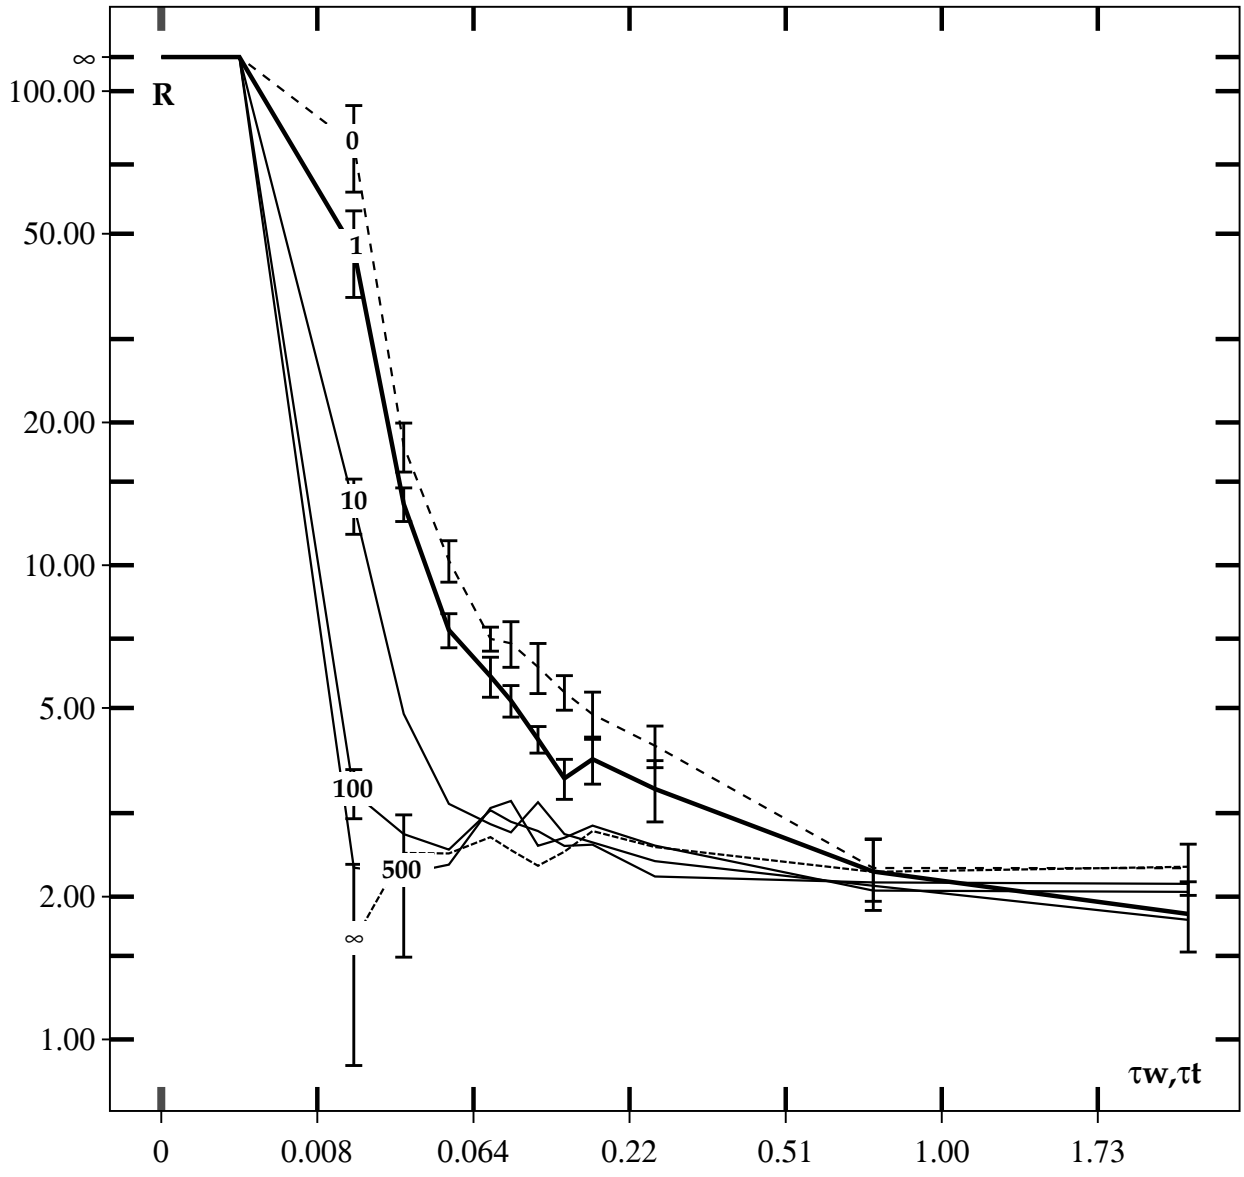
∞
100.00
R
0
50.00
1
20.00
10
10.00
5.00
100
500
2.00
∞
1.00
τw,τt
0
0.008
0.064
0.22
0.51
1.00
1.73

Figure 4:

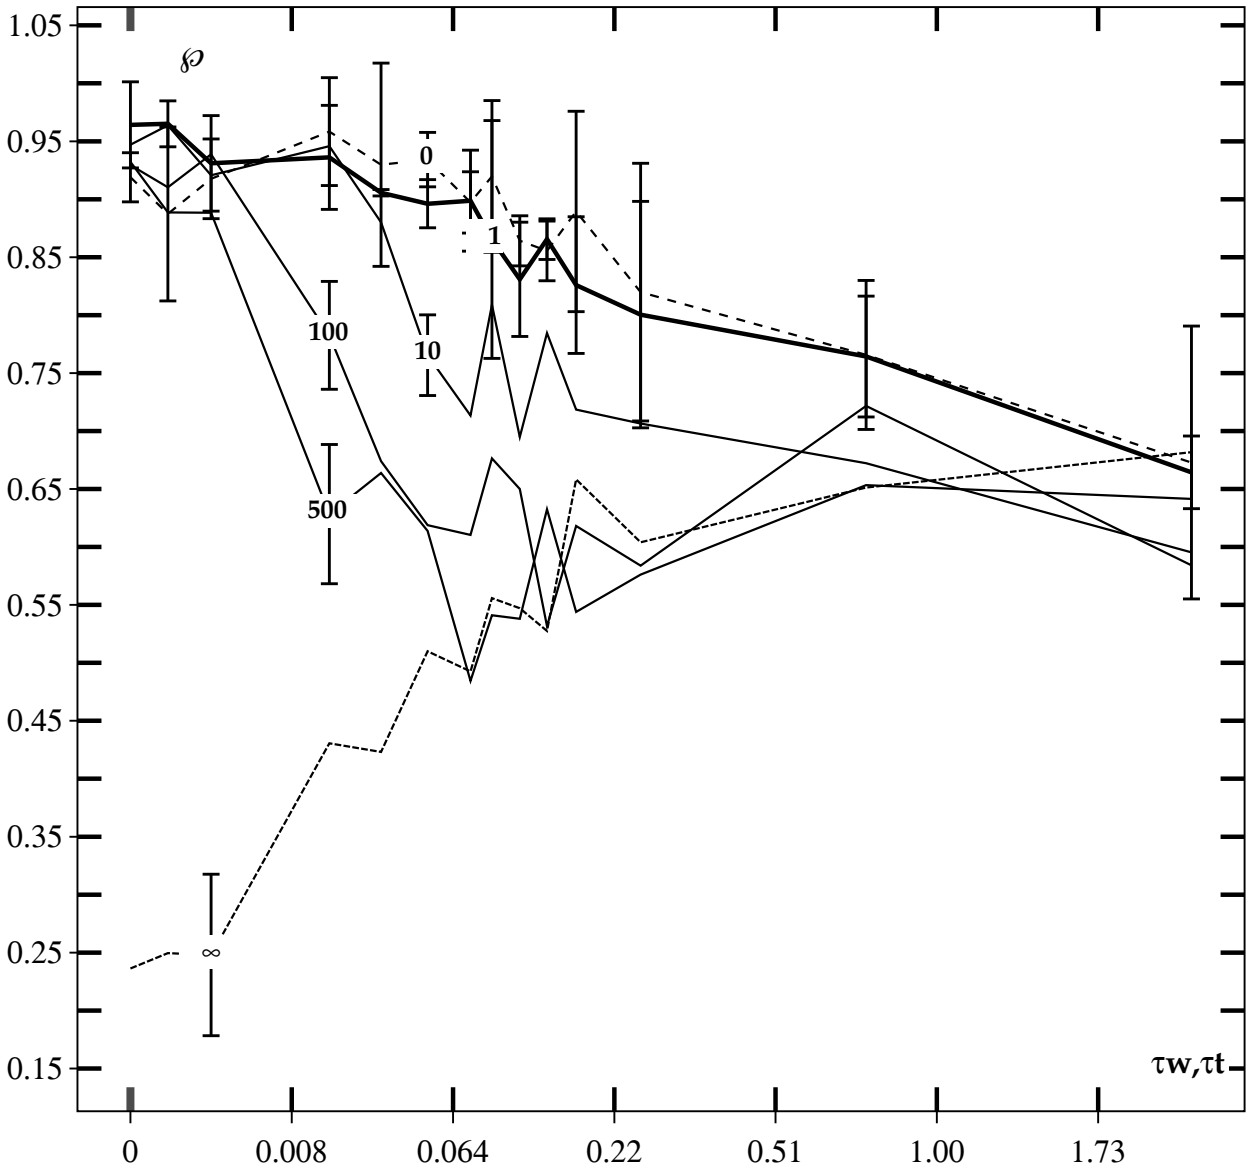
1.05
0.95
0.85
0.75
0.65
0.55
0.45
0.35
0.25
0.15
℘
0
1
100
10
500
∞
τw,τt
0
0.008
0.064
0.22
0.51
1.00
1.73

Figure 4b

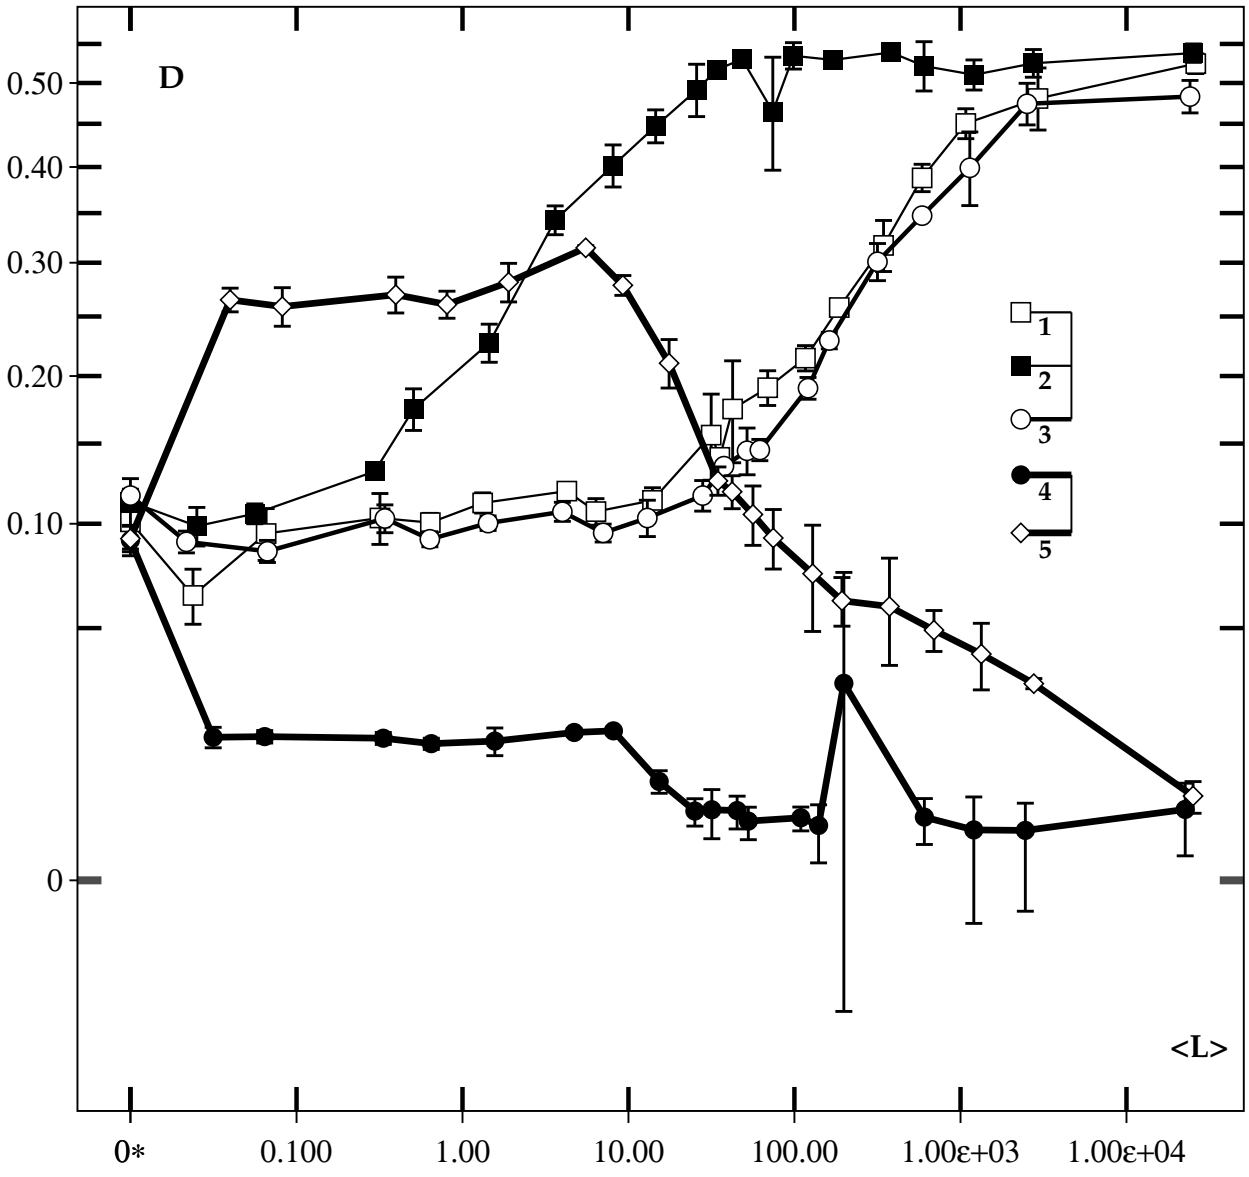
D
0.50
0.40
0.30
0.20
0.10
0
1
2
3
4
5
<L>
0*
0.100
1.00
10.00
100.00
1.00ε+03
1.00ε+04

Figure 5:

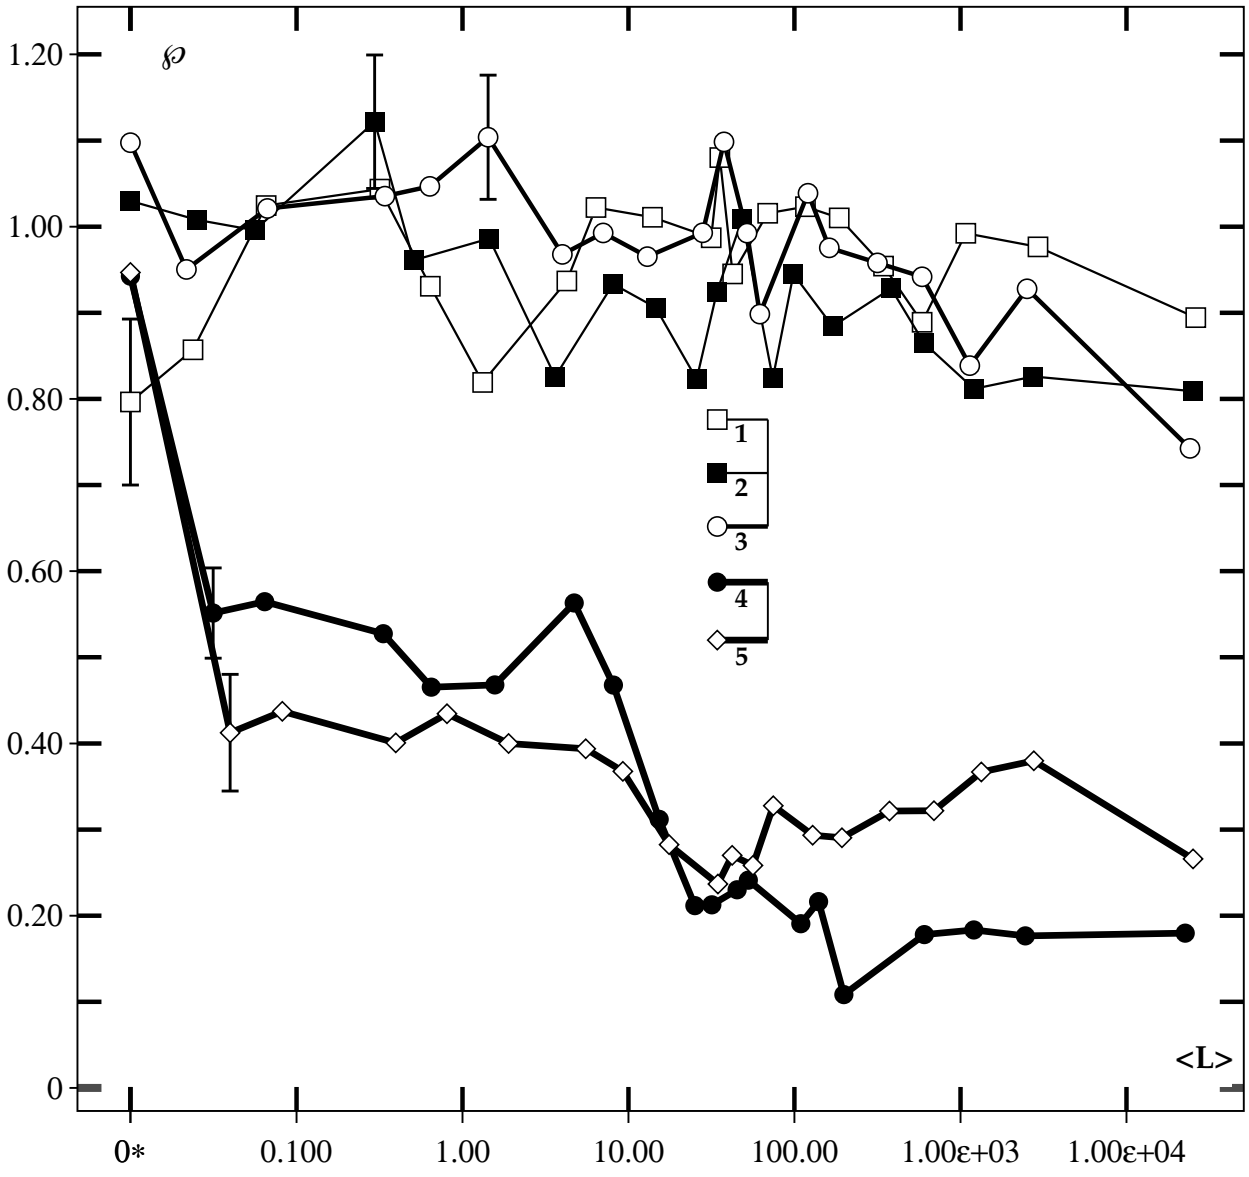
℘
1.20
1.00
0.80
0.60
0.40
0.20
0
0*
0.100
1.00
10.00
100.00
1.00ε+03
1.00ε+04
<L>
1
2
3
4
5

Figure 6: