Using Evolutionary Algorithms for the Optimization of the Sensorimotor Control in an Autonomous Agent

Susanne A. Huber
Max-Planck-Institut für biologische Kybernetik, Spemannstr. 38, 72076 Tübingen, huber@mpik-tueb.mpg.de

Hanspeter A. Mallot
Max-Planck-Institut für biologische Kybernetik, Spemannstr. 38, 72076 Tübingen, ham@mpik-tueb.mpg.de

Heinrich H. Bülthoff
Max-Planck-Institut für biologische Kybernetik, Spemannstr. 38, 72076 Tübingen, hhb@mpik-tueb.mpg.de

Abstract

Autonomous agents that evolve visually-guided control mechanisms using genetic algorithms (GA) and evolutionary strategies (ES) are introduced. These agents are situated in simulated environments and are designed based on the neurobiological principles of various aspects of insect navigation. They generate behavioral modules for obstacle avoidance and the compensation for rotations caused by external disturbances. The sensor positions and the visuomotor coupling evolve with the sensors and motors acting in a closed loop of perception and action. The probabilities for the genetic operations mutation and crossover are optimized and the results of the two optimization techniques genetic algorithms and evolutionary strategies are compared.

1 Introduction

In this work autonomous agents are introduced which navigate through a virtual world. Evolutionary algorithms – specifically genetic algorithms and evolutionary strategies – are applied to evolve visually guided control mechanisms which enables them to navigate through a virtual world. Behavioral modules for obstacle avoidance and the compensation for rotations caused by external disturbances are generated with the perception and action of the agent – sensor input and motor control – acting in a closed loop.

The agents are designed based on neurobiological principles of various aspects of insect navigation. The most important biological insight – applied to our artificial system – is that insects navigate mostly by evaluating visual motion information by means of neurons tuned to specific motion patterns (matched filters). The spatial localization of the receptive fields of these neurons are optimized with respect to certain behavioral tasks.

For genetic algorithms and evolutionary strategies the question of optimal parameter settings for the mutation and crossover probabilities is not yet fully understood. Therefore the probabilities for mutation (p_m) and crossover (p_c) are optimized empirically in this study. The results of the genetic algorithm and evolutionary strategies for various settings of p_m and p_c are compared.

In the following sections, results from the research on the visual system of flies are reviewed and the architecture of the autonomous agent is described. In section 4 the evolutionary algorithms used here are introduced, followed by the results of the simulations.

2 Perception of motion and visuomotor control in flies

For the visual orientation of flies the detection of motion plays a prominent role. While an insect navigates through a stationary environment the images on the retinae are continuously changing. This image flow depends on both the trajectory through the world and the structure of the environment (Longuet–Higgins & Prazdny, 1980). As a model for the perception of motion in insects, Reichardt and Hassenstein proposed a correlation detector (Hassenstein & Reichardt, 1956; Reichardt, 1961), where the temporal change of image intensities in the visual fields of two adjacent ommatidia are used.

From the investigation of the behavior of flies under controlled experimental conditions (Götz, 1964) as well as with freely flying animals (e. g. Wehrhahn, Poggio & Bülthoff, 1982; Wagner, 1985), different visuomotor subsystems e. g. course stabilisation, tracking and landing have been identified. The so–called tangential cells (Hausen, 1982; Hengstenberg, Krapp & Hengstenberg (in press)) in the lobular plate – a section of the visual system – are known to play a prominent role in the detection of egomotion and for the visuomotor control that compensates for deviations from the intended course. Position and size of the receptive fields of these neurons and also the specialization to certain motion patterns are essential for the course control of the fly.

3 The architecture of the agent

In order to build an artificial agent that navigates using strategies as they are known in flies, we attempt to evolve the matched filters and the sensorimotor coupling for certain behavioral tasks. In flies, each filter consists of a field of motion detectors with specialized orientations. In this paper we start with an agent that has only four visual sensors and two motors. Two sensors form a movement detector and the outputs of the two detectors are coupled via transmission weights to the two motors. The autonomous agent gathers information about its egomotion and the environment by evaluating the motion signals from the detectors. The orientations of the visual sensors determine which part of the motion field is used to navigate through the unknown environment. They are thus a particularly simple case of matched filters for the course control. Evolutionary algorithms are used to evolve the sensor orientations and the coupling weights between sensors and motors for different types of behaviors.

The visual signal of each sensor is the average of the incident intensities from 5×5 directions within the sensor's field of view. For each of the 25 directions – applying a ray-tracing operation – the intensity at the intersection of the line of sight with the visible surfaces is computed (Foley et al., 1987). The angular resolution of each sensor is $10°$ azimuth $\times$ $10°$ elevation. The time constants of the lowpass filters of the motion detectors are fixed ($\tau_1 = 2.0\text{s}, \tau_2 = 5.0\text{s}$). The matrix

$$\mathbf{W} = \begin{pmatrix} w_{LL} & w_{LR} \\ w_{RL} & w_{RR} \end{pmatrix} \tag{1}$$

contains the transmission weights for the sensorimotor coupling between the outputs d_L and d_R of the two motion–detectors and the two motors M_L and M_R. The velocity of the system is proportional to the force of the two motors, each motor producing a constant basic velocity $\mathbf{v}_0 = 5\text{cm/s}$ which is modulated by the motion detector outputs:

$$\mathbf{v} = \mathbf{v}_0 - \mathbf{Wd} \tag{2}$$

with $\mathbf{v} = (v_L, v_R)$ and $\mathbf{d} = (d_L, d_R)$. The temporal change of the rotation angle results from the difference of the velocities of the two motors

$$\dot{\varphi} = \frac{v_R - v_L}{c}, \tag{3}$$

where $c = 10\text{cm}$ is the distance between the two motors. The system has two degrees of freedom: rotation around the vertical axis and translation in the heading direction. The agents move through a tunnel with a width and height of 4m and a length of 100m. The height of the agent in the tunnel is kept constant at 2m.

4 Evolutionary algorithms

4.1 The genetic algorithm

In a simulated evolution – using genetic algorithms – the autonomous agents adapt to the environment. The positions of the sensors and the sensorimotor coupling of the agents are evolved. The free parameters of the system are (i) the viewing directions of the sensors which implicitly define the preferred motion vector of the detectors and (ii) the transmission weights. For both (i) and (ii) bilateral symmetry is assumed. These parameters are encoded in a sequence of Gray-coded bits – a bitstring. The positions of the optical axis of the sensors are described by the azimuth Φ and inclination Θ. Each angle is encoded with 4 bits in the range from $5°$ to $175°$ with a stepwidth of $11.3°$. The weights of the sensorimotor coupling are encoded with 3 bits each, with the decoded real values $[\pm 0.01, \pm 0.05, \pm 0.1, \pm 0.5]$. Hence 6 parameters forming a bitstring with a length of 22 bits are evolved.

From a random initial population of bitstrings the next generation is built, using the so–called roulette-wheel selection schema. Here the number of offspring of an individual is proportional to its fitness. Linear scaling of the fitness is applied before selection – with individuals of maximum fitness $f_{\max}$ being scaled to $n\bar{f}$ (n = scaling-factor, $\bar{f}$ average fitness of population). This prevents premature convergence of the optimization algorithm. In accordance with the so–called elitist strategy, the individual with the highest fitness is automatically transferred to the next generation (Davis, 1991). The selected parents exchange their genetic material by one–point crossover. As a second operator, point–mutation is used to introduce new genetic material into the population. In order to investigate how the optimization converges under variation of the scaling factor n, we run the simulations with $n = 1.2$ and $n = 2.0$. The population size is 100 for both conditions.

4.2 The evolutionary strategy

Here again the orientations of the sensors forming a motion detector and the transmission weights are evolved. For both the motion detectors and the transmission weights we assume bilateral symmetry. The parameters are directly coded as real numbers with the initial values for $\Theta, \Phi \in [5.0°, 175.0°]$ and $\mathbf{W} \in [-0.5, 0.5]$. The scaling factor here is $n = 2.0$, the population size is 100.

For the selection of the individuals that produce offspring for the next generation again the roulette-wheel selection schema and the elitist strategy are applied. The selected parents exchange their genetic material by one–point crossover. Mutation is introduced by multiplication of the parameter P with a random fac-

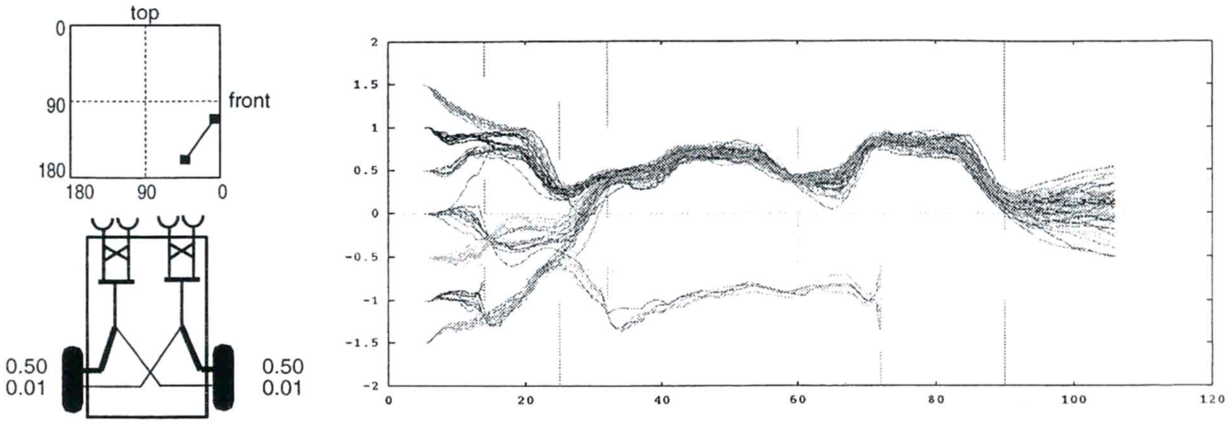

Figure 1: Sensor position on one hemisphere, weights for the sensorimotor coupling, and trails from different starting positions through a tunnel with obstacles. On sensory input and motor output $\pm 10\%$ noise is added.

tor $(1+r)$ where r is a random number, $r \in [-0.5, 0.5]$:

$$\tilde{P} = P(1+r) \; . \tag{4}$$

4.3 Probabilities for mutation and crossover

The evolutionary strategies of Rechenberg (1973) concentrate on mutation as genetic operator. The genetic algorithms of Holland (1975) use crossover as the primary operator with mutation being of secondary importance. The power of crossover lies in the construction of new and preservation of old individuals with high fitness. However the construction of new individuals becomes less effective as the population loses diversity, since the number of common alleles increases. In contrast mutation introduces new genetic material into the population and thus creates new individuals independent of the diversity of the population. According to Spears (1993) the choice of the genetic operators depends on whether the whole population should gain a high fitness – here using crossover is of advantage – or one optimal individual is to be found, in which case use of mutation is sufficient to obtain comparable and better results (Fogel & Atmar, 1990). In order to investigate the question of optimal parameter settings in our simulations, the probabilities for crossover and mutation (p_c and p_m) are varied (Tab. 1).

Table 1: Probabilities for crossover p_c and mutation p_m in 6 conditions.

	p_c	p_m
$C_{++}M_{+}$	0.70	0.01
$C_{+}M_{+}$	0.30	0.01
$C_{++}M_{0}$	0.70	–
$C_{0}M_{+}$	–	0.01
$C_{0}M_{++}$	–	0.05
$C_{0}M_{+++}$	–	0.10

5 Simulations

5.1 Obstacle avoidance

In these simulations the system has to avoid obstacles in the tunnel and maintain a safe distance of 10cm while passing by. The walls are at $x = 9.0$m and 27.0m, $0.0\text{m} \le y \le 2.0\text{m}$ and $x = 19.0$m and 40.0m, $-2.0\text{m} \le y \le 0.0\text{m}$. The texture on the walls, floor and ceiling is a sinusoidal pattern ($\lambda = 1\text{m}$). The fitness function used here is:

$$F = x(t_{\text{stop}}) \; , \tag{5}$$

where $x(t)$ is the actual position on the longitudinal axis of the tunnel and $t = t_{\text{stop}}$ is the number of timesteps the individual survived in the tunnel without colliding with walls. The maximum number of timesteps is 800.

We present here the results of three blocks of simulations. In each block the simulations are carried out with different settings for the mutation and crossover probabilities (Tab. 1). Each simulation is repeated 10 times. In block 1 and 2 we use genetic algorithms for the optimization but change the scaling factor from $n = 1.2$ to $n = 2.0$. In block 3 evolutionary strategies are applied using a scaling factor of $n = 2.0$. The initial population is constant throughout all trials. In all simulations agents evolve that navigate through the tunnel without colliding with obstacles.

In the example shown in Fig. 1 the motion detectors of the agent are diagonal with one sensor oriented in the direction of heading – perceiving obstacles the agent is approaching – and the other sensor oriented towards the floor. The velocity is

$$\mathbf{v} = \mathbf{v}_0 \tag{6}$$

and the temporal change in the rotation angle is

$$\dot{\varphi} = 0.51(d_L + d_R)/c \text{ deg/s} \; . \tag{7}$$

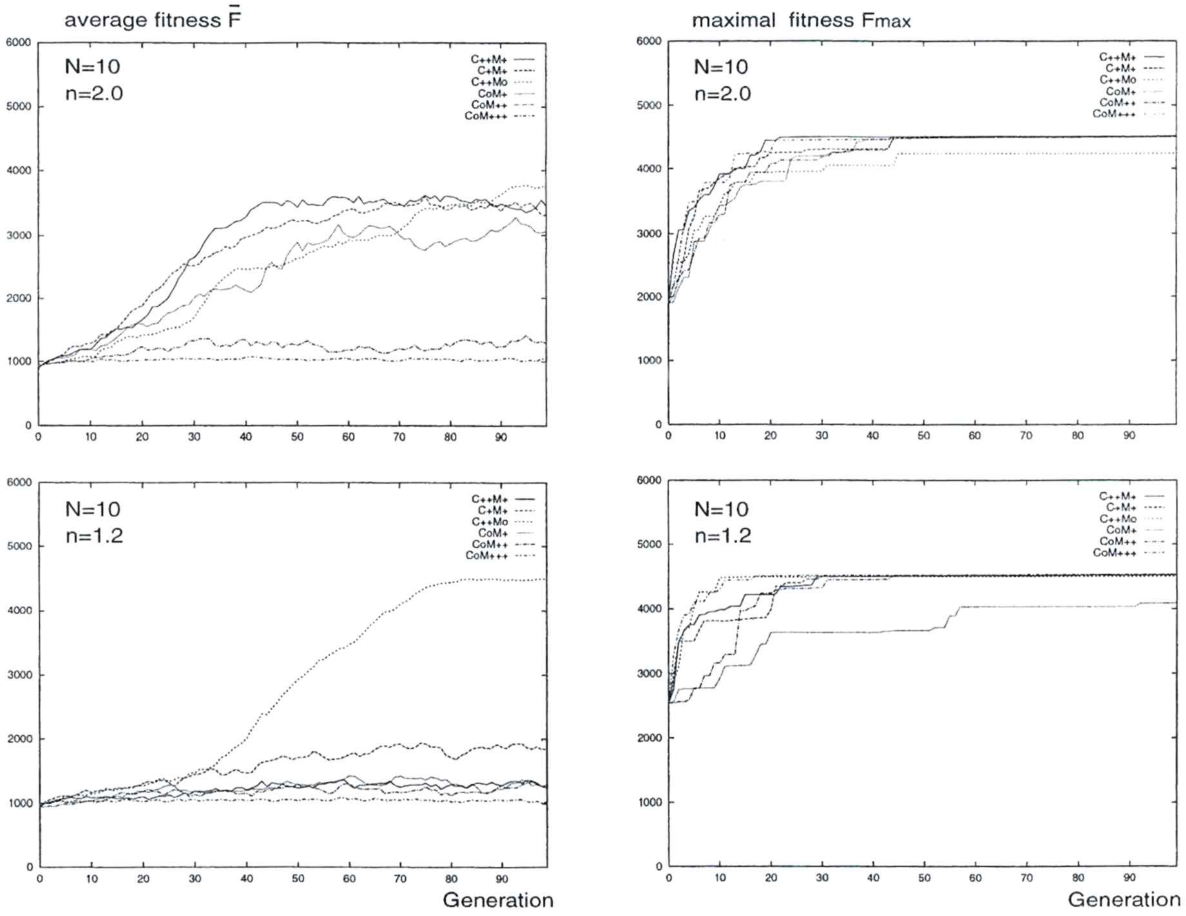

Figure 2: Genetic algorithms: Average (left) and maximal (right) fitness, averaged over 10 trials for different settings of mutation and crossover probabilities (see Tab. 1) with scaling factor $n = 2.0$ (top) and scaling factor $n = 1.2$ (bottom).

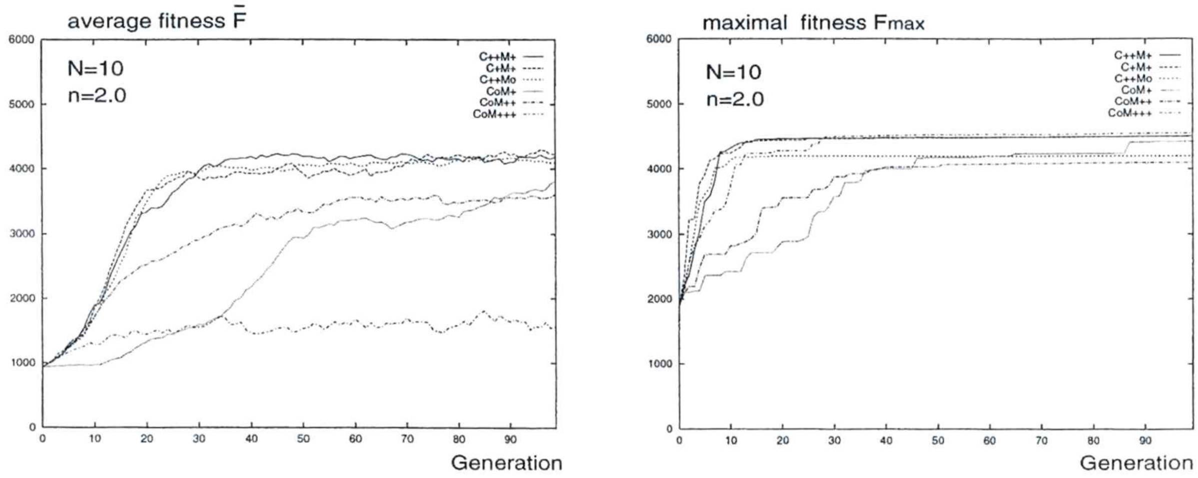

Figure 3: Evolutionary strategies: Average (left) and maximal (right) fitness, averaged over 10 trials for different settings of mutation and crossover probabilities (see Tab. 1) with scaling factor $n = 2.0$.

The architecture allows the agent to generalize the behavior to different environments. During the test phase we add noise of $\pm 10\%$ to the sensor input and motor output and vary the starting position of the agent (see Fig. 1). Under these conditions, it successfully navigates the tunnel in 98% of the trials for the starting positions $y = 0.0, 0.5, 1.0, 1.5$m and in 80% of the trials for all starting positions ($y = 0.0, \pm 0.5, \pm 1.0, \pm 1.5$m).

In Fig. 2 the average fitness $\bar{F}$ of the population and the fitness $F_{\max}$ of the best individual, both averaged over 10 trials are shown for genetic algorithms varying the scaling factor n. The mutation and crossover probabilities are varied according to Tab. 1. With a scaling factor of $n = 2.0$ (Fig. 2 (top)) the search for the individual with maximum fitness is fastest using a combination of crossover and mutation ($C_{++}M_{+}$). Using only crossover ($C_{++}M_0$) the maximum fitness converges to a suboptimum only. The average fitness $\bar{F}$ for C_0M_{++} and C_0M_{+++} show only a very small increase, indicating that with high mutation rates ($p_m = 0.1, 0.05$) the search is almost completely random. Crossover alone ($C_{++}M_0$), $C_{+}M_{+}$ and $C_{+}M_{++}$ lead to a high average fitness of the population.

With a smaller scaling factor of $n = 1.2$ (Fig. 2 (bottom)) the population hardly gains fitness except for the $C_{++}M_0$ and $C_{+}M_{+}$ conditions. However the search for the maximum fitness on average is speeded up, with $C_{++}M_0$ and C_0M_{+++} being the fastest, followed by $C_{++}M_{+}$ and $C_{+}M_{+}$. In this block all the conditions except C_0M_{+} ($p_m = 0.01$) lead to a high maximum fitness. This indicates that for the $C_{++}M_0$ condition one can prevent premature convergence using a low scaling factor, but for the C_0M_{+} condition $n = 1.2$ is not sufficient to gain a high maximum fitness.

The results of the simulations of the genetic algorithm in block 2 and the evolutionary strategies in block 3 (Fig. 2 (top) and Fig. 3) show no major differences in the optimization behavior, except that on average the convergence of the optimization procedure is slower in block 2 compared to block 3, resulting in a smaller average fitness at generation 100. In block 3 C_0M_{+++} shows the best optimization result, followed by $C_{++}M_{+}$ and $C_{+}M_{+}$. C_0M_{+} and C_0M_{++} show a very slow optimization behavior and converge together with $C_{++}M_0$ to suboptimal solutions only.

From these simulations, we can conclude that the most reliable optimization procedure is achieved with a combination of crossover and mutation. Here the population gains fitness (the search is not completely random) and an individual with high fitness can be found rapidly. Using only crossover and the scaling factor $n = 1.2$ seems to be another good strategy. However this does not lead to optimal results in general. An initial population with low diversity and individuals with low fitness will converge only to a suboptimal average fitness as no new genetic material is introduced using crossover only.

We are currently evolving agents in tunnels with a random dot pattern mapped onto the walls, floor and ceiling. First results of successful agents show that the sensors forming one detector have neighbouring positions and are oriented in the frontal area of the view sphere.

5.2 Compensation for deviations from the path

In these simulations the agent experiences a rotational force around the vertical axis. This can be caused by the slip of a wheel or by internal motor asymmetries. The agent has to evolve the parameters – using genetic algorithms – to compensate for the deviations and thus to avoid bumping into the walls. The same fitness function (eq. 5) is used. The maximum timestep is 1000. The rotational force has a sinusoidal modulation with a wavelength of 250 timesteps and a maximum causing a rotation of angular speed $\dot{\varphi} = 1.2$deg/s. Again the probabilities for mutation and crossover are varied for different simulations (Tab. 1), with each simulation being repeated 10 times. The scaling factor is $n = 1.2$. The results of section 5.1 are reproduced here. Again crossover alone leads to the highest average fitness followed by $C_{+}M_{+}$ and with combinations of mutation and crossover the individual with maximum fitness is found rapidly.

6 Summary and future work

Autonomous agents adapted to different behavioral tasks were generated using genetic algorithms and evolutionary strategies. The agents develop the viewing direction of their sensors and the sensorimotor-coupling in a closed loop and are thus able to compensate for deviations caused by external disturbances and to avoid obstacles. The influence of the crossover and mutation probabilities on the outcome of the simulations, concerning the maximum fitness and the convergence of the population was tested. Evolutionary strategies and genetic algorithms show a comparable optimization behavior. They achieve the most reliable optimization results with a combination of crossover and mutation. The individual with the highest fitness is found rapidly and the average fitness increases with the number of generations, indicating that the population converges.

In future work we plan to evolve agents which navigate through more complex environments. We will increase the number of movement detectors and use a linear array of sensors forming a 360° field of view. The agents will evaluate the motion detected in this field of view with filters that respond maximal to certain motion patterns – e.g. rotation around the vertical axis and translation in the direction of heading. In addition the agents will receive more degrees of freedom making 3D flight manoeuvers possible.

References

[1] L. Davis (ed.) (1991). Handbook of genetic algorithms. Van Nostrand Reinhold, New York.

[2] D. B. Fogel, J. W. Atmar (1990). Comparing genetic operators with Gaussian mutations in simulated evolutionary processes using linear systems. *Biological Cybernetics*, **63**: 111 -114.

[3] J. D. Foley, A. van Dam, S. K. Feiner, J. F. Hughes (1987). Computer Graphics – principles and practice, Addison–Wesley.

[4] K. G. Götz (1964). Optomotorische Untersuchung des visuellen Systems einiger Augenmutanten der Fruchtfliege *Drosophila. Kybernetik*, **2**: 77–92.

[5] B. Hassenstein, W. Reichardt (1956). Systemtheoretische Analyse der Zeit-, Reihenfolgen- und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers *Chlorophanus. Zeitschrift für Naturforschung*, **11b**: 513–524.

[6] K. Hausen (1982). Motion sensitive interneurons in the optomotor system of the fly. II. The horizontal cells: receptive field organization and response characteristics. *Biological Cybernetics*. **46**: 67–79.

[7] R. Hengstenberg, H. Krapp, B. Hengstenberg. Visual sensation of self-motions in the blowfly Calliphora. In: *Biocybernetics of Vision: Integrative and Cognitive Processes* (ed.) C. Taddei–Ferretti (in press).

[8] J. H. Holland (1975). Adaptation in natural and artificial systems. The University of Michigan Press, Ann Arbor.

[9] H. C. Longuet–Higgins, K. Prazdny (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society of London* **B 208**: 385–397.

[10] I. Rechenberg (1973). Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. Frommann–Holzboog, Stuttgart.

[11] W. Reichardt (1961). Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In *Sensory Communication*: 303–317, Hrsg: W. A. Rosenblith, The MIT Press and John Wiley and Sons, New York.

[12] W. M. Spears (1993). Crossover or Mutation? *Foundations of Genetic Algorithms*, Morgan Kaufmann, San Mateo, CA, **2**: 221–238.

[13] H. Wagner (1986). Flight performance and visual control of flight of the free–flying housfly (*Musca domestica L.*), III. Interactions between angular movement induced by wide– and smallfield stimuli. *Philosophical Transactions of the Royal Society of London*, **B 312**: 581–595.

[14] C. Wehrhahn, T. Poggio, H. Bülthoff (1982). Tracking and chasing in houseflies (*Musca*). *Biological Cybernetics*, **45**: 123–130.