Literature Review on Fly Vision
Final Report

Juha Kesseli, kesseli@cs.tut.fi

23.11.2000

Contents

1 Introduction

Considering the relative simplicity of its neural structure the fly has a very good vision. This can be frustratingly evident for people trying to catch one. For a couple of decades there has been research carried out on flies in order to learn the principles they use to sense the approaching ground, for instance. The response of flies to visual stimuli has been examined and analysis made.

For engineers the attractiveness of flies' vision lies above all in the simplicity mentioned above. The mammal brain is far too complicated to be presently understood at as detailed a level as the insects'. Still, even the capability to emulate insect vision would be an interesting aspect for developers of robots considering the performance of these systems. On the other hand, the fly vision can't be expected to do everything we might want from a visual system, since it's really a part of a larger system, the fly, that as a whole performs well in the chosen environment. The natural generalization of the topic would be animal vision as a whole. This document contains few references to other natural vision systems than that of the fly's.

This report contains a summary of the visual system of the fly in section two. The major features of photoreception and other physical parts of the system are included here. Especially in topics considering this part of the document it becomes evident that many parts of the fly's vision system are covered in such biological magazines that they are not easily accessed by engineers, and so the references to these kinds of articles are kept to a minimum. Pointers to some www-pages with bibliographies and research groups with references on the topic are given, however, in section 4. Many of the engineering articles have general fly information as references. Also especially the thesis work by Harrison has a lengthy list of publications.

Section three is devoted mostly to the engineering applications that have taken inspiration from the visual system of the fly. There are several research groups that have built VLSI implementations of motion detectors based on compound eye -type sensors. There has also been a lot of research concerning topics like object tracking required to chase potential mates in an environment full of obstacles or estimation of self-motion from the visual flow registered by the sensors. Some interesting systems have been designed and many principles that the flies use have really been shown to give simple and effective solutions to motion detection and robotic vision.

The articles that are used as references in this report are ones obtainable from the Internet either for free or with the online subscriptions of Tampere University of Technology. Overall, the collection of material mentioned has been made so that this report would give a good starting point for any further research on the topic. It seemed, however, sensible not to include all the references from the back of the Harrison's thesis, for example, in this document. That's why in order to find extra information in analog implementations the first thing to get after the references listed here is definately the thesis which is available for free on the web (see www-sources).

2 The visual system of the fly

Vision is a crucial instinct for flying insects. Over a half of the neurons in the fly's brain are believed to have something to do with the visual system. Perhaps the greatest difference compared with man-made systems is that the motion information processing is done at a local level with hierarchical information fusion. That is, the system shows a high degree of parallel processing [5]. The neural network in the fly's brain has the size of a pin-head and its weight is less than one milligram [2]. The different parts of fly's nervous system can be seen in figure 2.

2.1 The physical parts

2.1.1 The compound eyes

Each eye of the blowfly, for instance, consists of about 6000 lenses that project a single image onto the retina. For each of these lenses there are a total of eight cells responding to light stimulation. The basic unit of light sensing is called an ommatidium and it consists of one lens and the corresponding photoreceptors. Six of these receptors implement so called neural superposition, which pools together responses of the neighboring ommatidia in order to increase the effective lens diameter. The two other photoreceptors in an ommatidium are used only for sensing colors and are not involved in the most interesting aspect of vision, the motion detection. All in all, the response from a single ommatidium can be considered to represent the signal for a single pixel.

The incredible thing in the eyes is that the number of pixels is very limited. If the eyes were rectangular, the picture would be something like 77 x 77 pixels in a blowfly, which is very little compared with the systems commonly used in computer vision. In fruit fly the size is even less, 26 x 26 pixels in the array. In addition, with this number of pixels the fly sees nearly a complete hemisphere [5].

The first processing block in the system, the laminar region, also called the first optic ganglion, contains LMC (Large Monopolar) cells that amplify temporal changes in the data and suppress the DC level. The information is then passed on to the second optic ganglion.

2.1.2 The scanning retina

The retina in a fly's compound eye is not static. Instead, it makes scanning movements. These movements are able to increase the available resolution. The robustness properties of the following motion estimation are considered in [biologically inspired visual scanning]. The properties of the signals used for motion detection remain relatively invariant with respect to object distance and the grey-level of the object. That is, the time difference between signals of adjacent photoreceptors is not affected by these factors.

The scanning that Franceschini et al. have used in their robots is anterograde, that is the visual axes move in the posterior to anterior direction. If the translatory velocity and amplitude of the scan are known, in addition to the increased field of vision obtained, the distance to the contrast edges detected can also be computed. The scanning can be done in a number of ways that affect the properties achieved. The motion can be constant or variable [16]. See

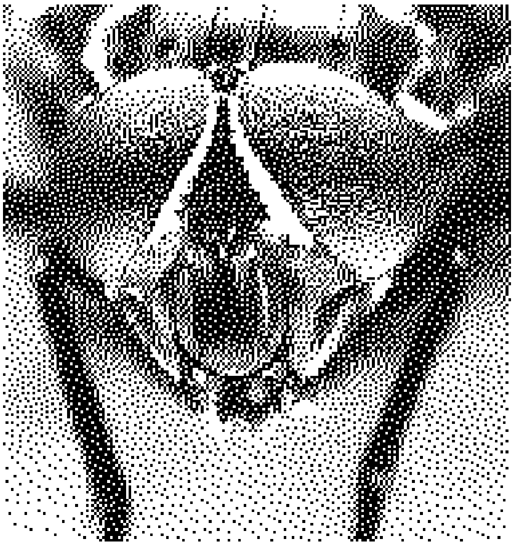

Figure 1: Fly's head and the compound eyes in a close-up shot [1].

also figure 8 and the text describing Franceschini systems in general. There are more references in the implementations-section.

2.1.3 The second optic ganglion

Very little is known about the structure of workings of the cells in this region because of their small size. It has been shown that the cells in the laminar region show highpass response that optimizes information transfer in this region.

Some recordings have been made from the cells in the medulla, and all kinds of responses have been observed, directional and non-directional motion responses included. The units estimating local motion in small parts of the visual field are called elementary motion detectors (EMD) or local motion detectors (LMD). The structure shows strong evidence for parallel processing, since under each ommatidium there is a similar structure of neurons.

2.1.4 The lobular plate

In this part of the brain the signals from all over the visual field are combined. The spatial convergence is quite remarkable and in the end there are about 50 tangential cells. Each of these cells transmits a signal that includes information from large parts of the previous stages.

Some of the cells respond only to horizontal or vertical movement. The 'H1'-neuron, for instance, responds to optic flow in front-to-back direction. The responses are not necessarily quite this simple, it has been shown that the H-cells (horizontal movement) can actually be modeled as matched filters corresponding to flow fields resulting from body rotations. This is discussed later in more detail.

There are also separate cells, FD neurons, that act as some kind of figure detectors. These cells respond mostly to small moving objects and are responsible for the incredible response times of male flies chasing potential mates, for instance.

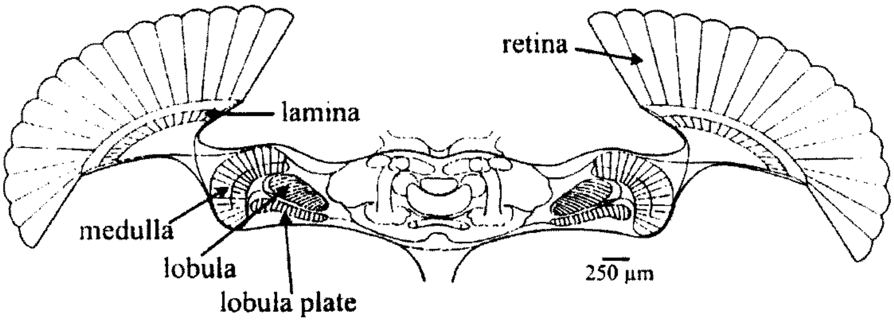

Figure 2: Over a half of fly's nervous system is presumably working in vision-related tasks [5].

2.1.5 Ocelli

Between the eyes, the flies have two low-resolution photosensitive organs called ocelli. The proposed functions of the ocelli are related mostly to horizon sensing. It has been observed that the ocelli seem to be involved in so called dorsal light response, which means that flies tend to align their head so that the center of brightness is straight on top. In outdoor environments this usually means that the head is kept straight up. Is seems also probable that the ocelli co-operate with the compound eyes on the task of positioning the head [5].

2.1.6 Polarized light detection

Fly has polarization-sensitive cells in its eyes . The ability to distinguish light polarity has been shown to be used by bees, for instance, as a compass system. This is possible since the reflected light is more or less polarized so the direction of the sun can be established even when it's cloudy. Fly's capability to take advantage of polarization has not yet been researched in detail [5].

2.2 Functional properties

2.2.1 Fast optomotor response

Flies have very short connections in their neural systems between sensors and actuators. The shortest reaction times reported have been around 30 ms when a male fly was chasing a prospective mate. The information does not have to pass through all the levels of the brain in order to generate a response to the changes in the surroundings. Similarly, the stability during flight is preserved by automatically generating needed torques with the wings when excessive problematic movement is observed in the optical flows.

The flies can extend their legs only 70 ms before the landing. The estimates for the landing moment are probably generated by integrating in time certain tangential cells that respond to expanding optic flows. Similar techniques are probably related to the escape response which is triggered by both motion and the decrease in average light intensity.

2.2.2 Stabilizing gaze

It has been shown that flies use head rotations to keep their gaze stabilized for most of the time also in banked turns, for instance. Their head moves much less than their thorax with respect to their environment except for short saccade movements that are just about as fast ast the photoreceptors' integration time.

When the saccade movement is not occurring, the fly can gather information from its surroundings with relative ease as the rotational component of visual flow has been mostly eliminated. Although it is in theory possible to separate rotational and translatory visual flows, it's hard in practise with limited performance of the components of vision. Overall, the stability of the head is so good that it's not possible to explain it only with vision. Other elements, such as the so called vestibular sense, a kind of gyroscopic system, are needed as well [6].

2.3 Related sensor systems

In many of the tasks that vision is mostly concerned with there are also other senses in action. Probably the most important of these is the vestibular sense. Real flies have a kind of angular rate gyroscopes attached to their body. With these so called halteres that beat up and down at the wingbeat frwquency the fly can sense angular rotation of the body around all three axes. This is achieved by integrating Coriolis force effects over a period of motion. This information is used in association with sight for controlling flight. With both halteres removed a fly quickly falls to the ground. A CalTech-based research group has a project on building a micromachined gyroscope based on the working principle of the halteres (see WWW sources).

It is also presumed that flies can detect linear acceleration by sensing the inertia of their head and limbs. There are cells in the neck and legs that measure both position and strain that could be used for this purpose [5].

3 Principles of Fly's Visual Information Processing

As of yet we don't really have the capability to internally examine all aspects of the signal processing in a fly. This means that a large part of the research has to be carried out by analyzing the behaviour of flies. Man-made systems that are based on insect vision therefore usually use some of the probable insect techniques by applying models that correspond to the external behaviour of flies. The models are then implemented as robotic systems, for example. The models usually have quite a good degree of experimental evidence from nature to back them up, however.

There are numerous systems built this way and the most interesting concept imitated from the flies is the motion detection using LMDs (local motion detectors). This is why models of motion detection are discussed first.

3.1 Visual information processing in general

Flies use several simple mechanisms for perception. For instance, they don't have a complex 3-d preception of the world, instead they mostly cope with simple measurements of apparent velocities of objects' images. The elementary motion detectors measure the angular velocity of an edge passing by a point in the retina. This information can then be used for localizing small objects to follow or to sense the self-movement with respect to the environment by integrating the responses of all the EMDs, for example.

These and other known aspects of fly's visual information processing system are presented next.

3.2 Motion detection principles

The most used aspect of fly vision in applications is the use of elementary motion detectors that enable the processing of motion information in parallel structures. The extraction of local motion information can be done in several ways [5].

3.2.1 Feature-based motion detection

Feature-based motion detectors look for specific features in the image seen. Typically, these look for changes in spatial or temporal features.

The spatial feature method is not that local at all, as it looks for objects found in an image from the previous image. The corresponding locations tell the speed that the object was moving at. This procedure requires some sort of segmentation before anything can be done and because of its character it's not that local at all. It's not very probable that flies use this method at least in the most significant tasks.

The temporal feature method looks for changes in intensity at pixels and compares these with known patterns emerging when an object is moved in the visual field. These known patterns are called templates by the Bugeye research team that uses them in their insect-vision robot [21]. This robot has achieved quite good results with template models.

3.2.2 Correlation-based motion detection

The principle described here corresponds perhaps the most closely to the observations made on flies and is used, for instance, in the Harrison's device [5]. Nicolas Franceschini's research group employs the same technique.

Elementary motion detectors are located in the medulla layer of the fly's optic lobe. These cells are small and hard to record from, so not much is known about their inner workings. The unit's outputs are, however, modelled very well for application purposes. The unit as described here is called a Reichardt motion detector.

3.2.3 Reichardt Motion Detector principle

The two basic properties of EMDs are the geometric constraint and the kinematic constraint. The geometric constraint states that the range of vision decreases when the direction considered gets closer to the line of self-movement. The kinematic constraint states that the range of vision in any particular direction increases with increasing self-motion speed (in case of stationary targets) [16].

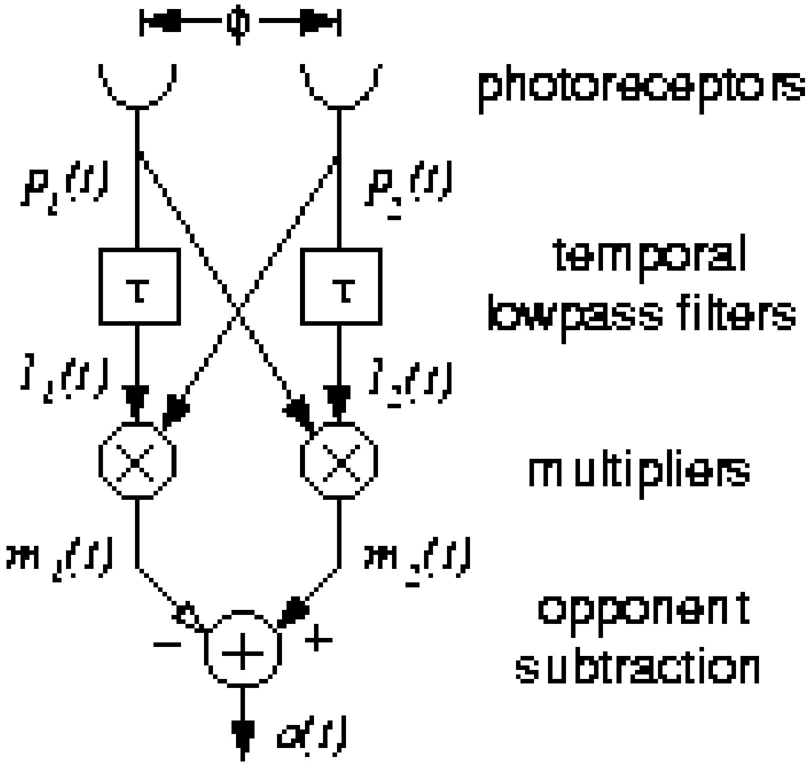

Figure 3: A correlation-based elementary/local motion detector [5].

The model commonly used to describe EMD behaviour is based on correlations between adjacent photoreceptors (figure 3). The incoming signal is first bandpass filtered to remove the high-frequency noise components and the low-frequency information (constant intensity levels) which is of no interest in motion detection. The edge information is all that is interesting for EMD purposes. The signal formed in this way is then multiplied by the corresponding, but delayed, version of a signal received from a neighboring photoreceptor. In the picture a lowpass filter is shown to represent an implementation of this delay. An analysis of Reichardt motion detector is available [5, pg 16].

The subtraction in the figure takes into account the fact that one block as described performs motion detection only in one direction. After the subtraction we have a signed output with the sign indicating direction of motion. This figure is taken from an engineering article, and represents a common way of forming an EMD output. In these applications it's normal to have a single row or at most a few rows of photoreceptors in the sensor unit. This is why all that's really needed is an amount of EMDs indicated, two for each pixel.

The bandpass nature of the signals used has been observed also in practise. The consequence of this is that the amplitude of the response varies with speed of the detected edge. Interestingly, this variation has been of use in the man-made systems as well as in flies. The effect is to prevent overcompensation oscillations and to enable a smooth movement. The passband is centered around the frequencies corresponding to typical movement speeds in a fly's environment.

3.3 Implementations

3.3.1 General

In a real fly the photoreceptors form a 2-d pattern on the retina and there must be connections in vertical direction as well, so the total number of multipliers in a 2-d image application would be four. In the real fly the motion in each direction must be processed somehow from about 50000 photoreceptor inputs. The applications of EMDs built have had, for instance, six rows of 24 pixels in the recent CalTech system has [5]. The Bugeye-group's device has a single 61-pixel row [21].

Most of the implementations mentioned here are built on wheels and are not anywhere near fly-size ([16] system has a base diameter of 30 cm). In addition, in laboratory conditions, there is usually a PC connected to the devices at least for diagnostic purposes. A notable exception is the Micromachined Flying Insect project at Berkeley University. They have been mostly developing the mechanical aspects of a flying insect, like the wings and the thorax. They have found that piezoelectric actuators and flexible thorax structures can provide enough power density to the wings. It's also fascinating that solar cells can supply adequate power. The micromachined gyroscope built in CalTech is also related to this project. The goal of this project is to develop a fully autonomous fly-sized robot (see www-sources). These projects have not yet been significantly involved in the vision research.

Some general advantages of using analog processing (so called neuromorphic chips) in early vision are given by Koch in article [11]. The main point of the article is that it's unnecessary to do complicated smoothing operations in the digital domain, if an analog neuromorphic vision chip can do the same thing with simple operations. There are lots of publications on neuromorphic vision chips in general (See the www-addresses). More powerful systems including also digital parts are considered in [7], for example. Further information on these systems can be found from the references cited in these articles. This is not strictly fly-related, however.

The most common implementations of neuromorphic chips are the so called silicon retinas that could be used to replace CCDs in applications like pattern recognition, if the resolution needed is not too high. The main advantages of silicon retinas are that they can adapt to a wider range of relative light intensities and also the smoothing of the image is done in the chip, so no power-hungry computations are needed in the digital system attached [11]. The main advance made in recent years in CalTech in the sensor area has been that the outputs of the bandpass sensors have become quite well independent of the brightness levels and also distances involved. This has not been the case in the early analog sensor times [11].

3.3.2 Modeling of the fly's compound eye

The research group led by Moya intends to build a comprehensive model of the fly's eye at a very detailed level. There is work done already, for instance there is a model of the photoreceptor cells at the level of ion flows occurring in each of them [15]. This model has been observed to perform in a manner similar to that of the real fly. Moya's dissertation, published in 1997 at University of New Mexico, was entitled 'A comprehensive model of the neuro-ommatidial layer in Musca'. Other aspects already taken into account in the model include electrocoupling affecting the photoreceptors and the adaptational properties of the photoreceptor membrane. A hardware model of the system is being considered.

An analog implementation of the fly's photoreception model has been built by Wilcox and Thelen [19]. The article describes several factors affecting the resolution achievable and the phenomenon typical for animals called hyperacuity. This means that the resolution achieved is actually better than the coarse array of photoreceptors would suggest. Their construction demonstrates hyperacuity in machine vision for the first time.

Compared with the silicon retinas mentioned earlier, these devices use the same basic principle but also take into account the properties specific to flies.

3.3.3 Small-sized scanning visual sensors

Whereas most applications discussed use constructions not that small at all, at least a Japanese research group has taken interest in reducing the size of the visual sensors. The latest available information describes an artificial eye with a radius of 30 mm. This eye has the scanning motion feature implemented and it has a total of 60 photoreceptor elements. The team's goal is to approach the size of a real fly's compound eye with the radius of about 500 um [9].

Similar sensors have been studied also in CalTech, and in the www-sources section there is a link to research report entitled 'Visual Sensor with Resolution Enhancement by Mechanical Vibrations'. The report discusses the mechanism by which the resolution enhancement can be done on the neural signal level by interpreting signal spikes transmitted by the scanning sensors. This research does not refer specifically to flies.

3.3.4 Template-based Bug-Eye

The Bug-Eye is the robot of the Australian research group including Yakovleff and others. Their Bugeye I had the mentioned 61 pixels in a single row. The surprising thing is that with this simple setting a robot moving on the ground has been able to move relatively well in a test environment and the systems have been proposed for use in obstacle avoidance applications with requirements of low power consumption and light weight [21]. The power consumption of Harrison's EMD-array -based system is less than 5uW. This is a great achievement. For comparison purposes, traditional CCD imagers used for similar purposes on Sojourner rovers in recent Pathfinder Mars missions used 0.75W without any data processing.

Figure 4 shows some data recorded from a Bugeye. The letter codes correspond to intensity changes with certain polarity and magnitude. This data is the starting point for the analysis that follows. In order to recognize edges that belong to the same object, a tracking algorithm is used. It keeps track of the edges by keeping an eye on them, and recognizing looming patterns, for instance. This kind of a pattern is a clear sign to turn quickly before a collision occurs. Flies use this kind of simple information too, no information on distances or self-motion is needed to come to the obvious conclusion.

It's also possible to measure distances to objects that are assumed stationary and compute estimated times before contact is made with an object. These kinds of calculations are shown in figure 5. The distance calculation assumes knowledge of self-motion velocity, but this is usually available in ground-based robots. Self-motion estimation from visual information is discussed elsewhere in this document. Flies use integration over visual field and air speed measurements to measure self-motion so this information is available in real flies as well.

The Bugeye is constructed from analog VLSI circuits, and an example of a circuit used is shown in figure 6. In order to relate different edges to one another, the Bugeye uses velocity matching. The two voltages driving this circuit are obtained from other circuits that give the intensity change speed at an edge. A voltage-buildup mechanism is used to indicate similar velocities for two edges. This kind of techniques are, of course, not easy to use in a general environment. For instance, a gap between two blocks must be distinguished from a block between two openings. Some objects are bright on a dark background and some have it the other way around. An important clue is given by knowledge of the type: “Objects are typically made of material with constant brightness”.

Other area where different techniques are used here is the edge tracking. The systems tested usually keep track of any edge-like patterns detected and make conclusions after the track has been active and tracked for several rows.

3.3.5 Correlation-based systems

Robots based on previously presented correlation principle include the one built by Harrison et al and others developed by the long-term contributor Franceschini with his research team. These follow maybe closer the actual processing of information in a real fly. The Franceschini system has been selected as the primary topic of discussion because it has many interesting, yet simple to consider, features. For further information on other implementations, see the literature

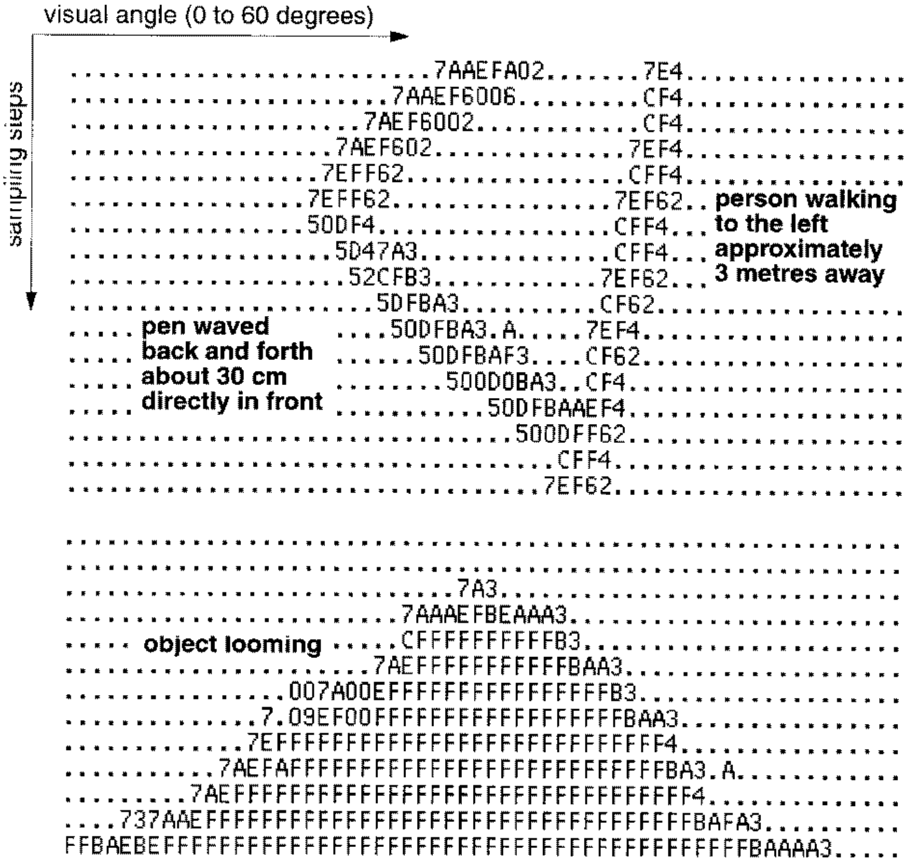

```
..........................7AAEFA02.......7E4................
.......................7AAEF6006.........CF4................
.....................7AEF6002............CF4................
...................7AEF602..............7EF4................
..................7EFF62................CFF4................
.................7EFF62................7EF62.. person walking
.................50DF4.................CFF4... to the left
...................5D47A3..............CFF4... approximately
....................52CFB3............7EF62... 3 metres away
......................5DFBA3..........CF62..................
..... pen waved         ....50DFBA3.A.....7EF4...................
..... back and forth    ......50DFBAF3....CF62...................
..... about 30 cm       ........500D0BA3..CF4....................
..... directly in front ...........50DFBAAEF4....................
................................500DFF62....................
...................................CFF4.....................
..................................7EF62.....................

............................................................
............................................................
............................7A3.............................
........................7AAAEFBEAAA3........................
..... object looming .....CFFFFFFFFFFB3.......................
......................7AEFFFFFFFFFFFBAA3....................
................007A00EFFFFFFFFFFFFFFFFB3...................
..............7.09EF00FFFFFFFFFFFFFFFFFFBAA3................
.............7EFFFFFFFFFFFFFFFFFFFFFFFFFFFF4................
...........7AEFAFFFFFFFFFFFFFFFFFFFFFFFFFFFFBA3.A...........
.........7AEFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF4..............
....737AAEFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFBAFA3..........
FFBAEBEFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFBAAAA3.....
```

Figure 4: Some intensity changes detected by the Bugeye with descriptions of the sources [21].

([5] at least).

There are many versions of Franceschini's robot but the basic principles are about the same. This version is in accordance with [13]. The movement of the robot is divided to two phases. During translatory movement in straight line at constant speed the LMDs do their job and calculate local motion estimates. After the straight line movement a rotation is made based on the motion detected and a new velocity is selected for the line segment. During the rotation the robot doesn't move forward at all. This simple principle has avoided the problems of separating translatory from rotatory movement. The robot described here has two eyes with 15 optical axes each. These move apart from each other to a distance comparable to the velocity. This is done in order to avoid a dead zone straight ahead.

The correlation-based Local Motion Detectors drive so called EBCs (Elementary Behavioral Circuits). There are several of these just like LMDs. Each EBC makes its own proposal as the new velocity and direction of motion. Velocity is proposed to be such that the target causing the contrast will remain tangent to the circle of vision during the movement. The EBC generates two direction outputs, plus and minus something close to 90 degrees from the target

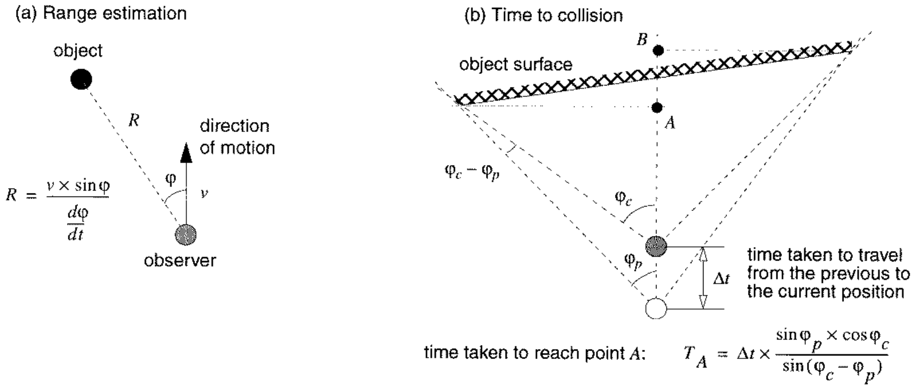

Figure 5: Simple formulas that can be used to obtain range estimates in both distance and time to collision [21].

if it's an object to be avoided or turn towards the motion if it's a target to be chased.

The velocity affects the circle of vision as mentioned previously, so the selection of velocity EBC proposes is such that the robots sort of keeps an eye on the object during the translation. Flies have been observed to fly at such a ground speed that the retinal slip stays constant, so this principle is very closely the same used by real flies. The rotation selected results in repulsion of obstacles.

The speed of the entire system is selected to be the minimum of those proposed by all the EBCs and a target detector that tries to increase the speed every time an object to be chased is in view. An important point to keep in mind is, however, that the velocity must not be too slow, or the robot might drift too close to an obstacle resulting in a collision. The steering angles proposed towards the right and left are averaged separately and then the one of these requiring a smaller rotation by the robot is selected as the final result. This strategy results in excellent performance in a forest of obstacles shown in figure 7. The robot avoids obstacles by keeping them at a radius corresponding to the circle of vision.

The exciting thing about this construction is that it shows that a great deal of the information processing can indeed be done in parallel and the resulting elementary movements can sensibly be then combined to a single command. The system does not have an inhibitory/suppressive system as many decision making robot guidance systems do. All the elementary decisions contribute in a cooperative and competitive setting. This kind of a system seems to be very capable of doing low-level, reflexive type of movement control. A higher-level planner could be built on top of this to create long-term strategies and define targets.

The same research team has also included a scanning retina -type of a system in their robots and positive results of these experiments can be found in [16], for example. Figure 8 shows the improved accuracy obtained. The gray areas in the pictures show the size of an area in which objects cannot be separated for ordinary and for a scanning retina.

The system discussed in Harrison's thesis contains a description of an interesting implementation. As mentioned later, as its greatest achievement it implements a tight optomotor response loop. In essence, the yaw motion estimated from the images is used as negative feedback for the guidance system giving orders to turn. The system didn't always succeed as well as it did most of the time, and this was accounted to the fact that there were areas in the office that had quite a constant colored background.

The sensors are tested for robustness and direct comparisons with the fly are made. It is found that the system performs in many ways like a real fly in its direction stabilization task. The wide-field motion-sensitive neurons were also modeled in great detail, and this improved the performance by decreasing the effect of patternless areas on the final result (see self-motion estimation).

3.4 Optomotor response, self-motion estimation

The so-called optomotor response of the fly is one of the most interesting engineering applications. It means the ability of the fly to move to the desired direction despite the effect of winds and the imprecision in the flight mechanics. This kind of a task is relatively difficult to implement with traditional approaches.

The basis of optomotor response is formed by low-level connections between sensors and actuators, which enable automatic corrections to individual wing speeds and angles when necessary. This kind of effects have been studied in depth. It's currently possible to produce correct behaviour of the fly's optomotor response by feeding computer-generated data directly to a fly's brain bypassing its eyes.

This kind of a system is also representative of the capabilities of several robots, especially of that built in CalTech [5]. The vehicle built had five times larger gear ratio on the other side. As a result, without any guidance it kept moving in small circles. With the optomotor response loop constructed the vehicle was able to move in a straight line extremely well.

A significant improvement was also made by imitating the performance of flies' wide-field motion-sensitive neurons more closely (these are the integrating, tangential cells). A nonlinear behaviour modeled at the ion flow level is taken into account. If the responses of LMDs passed on to integration are conveyed linearly, the size of the pattern on the retina affects the performance (the amplitude of the response). The nonlinearity enables the integration to perform rather well despite large areas on the retina where there is little or no movement. This is implemented in practice with varying conductance components and the effect of the gain control on power consumption was very small [5].

In the thesis it was also proposed that local spatial derivatives could be used to estimate the reliability of the vision-based information. This could be important when, for example, sensor fusion is performed with visual and gyroscope sensors.

There are also several papers on analog velocity estimation sensors, for example [10]. These are basically doing motion detection in the way that is done in the fly applications, using simple analog circuits. There are, however, several different techniques used like zero-crossing detection and gradient techniques, that do not perform the way flies probably do. These are really just a part of the category of analog neuromorphic vision chips, and are also mentioned in [11].

It's been proposed that at least bees (why not flies too) use motion parallax to perform segmentation on visual scenes [8]. The analog velocity sensors mentioned can be used to create VLSI motion discontinuity detectors. These have been tested and shown to have good performance [12]. Considering the nature of their operating principle, they could indeed be the kind of systems operating in flies. No knowledge of the flies' segmentation principles was found, however, so this just goes to show what kind of things neural equivalents corresponding to these analog devices could in principle do. The complexity of the comparable neural systems is hard to judge, and many of these solutions might model systems in mammals at most, and not flies, for example.

The references in [12] include a book called 'Analog VLSI and Neural Systems' that probably has information related to the analog circuits discussed here. The articles by Koch and others also include references in the theme.

3.4.1 Matched filters as self-motion estimators

The response of the integrating tangential neurons has been investigated in detail, and it has been observed that the local motion sensitivities and directional responses are not uniform. Instead, they bear striking resemblance to flow fields that are induced by rotations of the fly around different axes, for instance.

As a consequence of this, the response of a tangential neuron to a local motion has been modeled as a projection of the movement on the preferred motion direction multiplied by the local motion sensitivity. These local responses are then added to create the neuron output. This model is, in effect, a matched filter. Matched filters have been in practise used for self-motion estimation [4].

The self-motion parameters can't be deduced from the flow-field alone, since the local speeds depend on the distances to objects at corresponding directions. This is why, in general, the optimal weights for the matched filters must be generated by using statistical information of the surroundings. The basic statistics that have been used for a robot in office environment are the distance means and standard deviations in all directions [4].

Different sorts of analog velocity detection schemes are discussed in articles by Koch and others, for instance (the CalTech research group).

4 WWW sources

The table 1 contains several links to different sites on the web that could be of interest for anyone studying fly vision. The most interesting ones are briefly commented here.

Bug-EYE group is a research team at the University of Adelaide specialized on creating algorithms explaining insect vision and building analog VLSI implementations of visual systems inspired by insects. They have built several versions of their chip called Bugeye. The papers concerning their work, including [20] and [21] can be downloaded directly from this site. Their implementation uses templates (temporal feature motion detection).

The Laboratory of Neurobiology in Marseille has a research team lead by Franceschini that has done a lot of research on the topic. The pages include listings of their publications, but many of them are in French and are probably hard to access. Fortunately, the major papers have been written in English as well and are easily accessible, e. g. [2], [13], [16], [18]. Other interesting, but probably harder to obtain, material includes a thesis work by Mura (As a reference in [9]).

Reid Harrison's Ph. D. thesis [5] was published in May 2000 in California Institute of Technology, and it is available from his homepage at the University of Utah. The topic of the thesis was an implementation of a VLSI motion detection (with a correlation-based detector) and analysis circuit. The thesis contains an excellent bibliography on almost any topic regarding fly vision including the biological articles, most of which are not, unfortunately, accessible with the University subscriptions. The key points of all aspects of fly's sensor system are summarized in the early chapters and also along the way as the system's properties are discussed. The bibliography contains also engineering articles that have not been listed in this document's references.

The MFI and Biomimetics Robotics research groups have mostly been involved in developing mechanical parts needed to build insect-like robots. The sites have lots of interesting material. The Biomimetics project in Stanford is concentrated on creating walking robots. These are also studied in many other places around the world.

The MFI site contains information on the interesting Micromachined Flying Insect project, where the aim is to build a flying insect-sized autonomous robot. The site also contains information on the performance characteristics of real flies. The project intends to integrate simple optical sensing on the robot, but no information concerning this aspect of the work was found.

The Dickinson's lab has publication lists on fly aerodynamics and also several video clips.

Hans van Hateren's labs vision group has investigated fly's retina and lamina as well as motion detection in flying insects. The pages contain links to articles but also images demonstrating the fly's poor visual capability in the dark, for instance. There are also video clips showing fly movement in flight. These are related to studying the fly's capability to keep the eye relatively well in the straight position in banked turns. The most interesting video clips are maybe ones that show what a fly would see looking at a room. These clips have been created using a theory that explains preprocessing of images based on maximization of information, given normal properties of natural images, and basic properties of fly neurons. The results also clearly show that flies cannot see that well at all when the level of illumination is decreased.

References

[1] S. Bains, Providing robots with a fly's eye view, SPIE OE reports, `http://www.spie.org/web/oer/september/sep99/cover2.html`, 1999.

[2] N. Franceschini, J-M. Pichon, C. Blanes, Real time visuomotor control: from flies to robots, Fifth International Conference on Advanced Robotics, 91 ICAR, 1991.

[3] N. Franceschini, Vision, flies, neurons and robots (abstract only), Proc. of 1995 IEEE International Conference on Robotics and Automation, IEEE, 1995.

[4] M. Franz, T.R. Neumann, M. Plagge, H. A. Mallot, A. Zell, Can fly tangential neurons be used to estimate self-motion?, Artificial Neural Networks IEE Conference, 7-10 September 1999, 1999.

[5] R. R. Harrison, An Analog VLSI Motion Sensor Based on the Fly Visual System, Ph. D. thesis, California Institute of Technology, 2000.

[6] J. H. van Hateren, C. Schilstra, Blowfly flight and optic flow. II. Head movements during flight, J. Exp. Biol. 202:1491-1500 (1999), 1999. (preprint at `http://hlab.phys.rug.nl/papers/phead.html`)

[7] C. M. Higgins, C. Koch, Multi-chip neuromorphic motion processing, 20th Anniversary Conference on Advanced Research in VLSI (IEEE), 1999.

[8] G. A. Horridge, What Can Engineers Learn from Insect Vision?, Conference Proceedings, 1993 International Conference on Systems, Man and Cybernetics, IEEE, 1993.

[9] K. Hoshino, F. Mura, I. Shimoyama, A micro-sized visual sensor based on the fly's compound eye with a scanning retina, Technical Digest, 12th IEEE International MEMS Conference, 1999.

[10] G. Indiveri, J. Kramer, C. Koch, System Implementations of analog VLSI velocity sensors, IEEE Micro, vol. 16, issue 5, pg. 40-49, 1996.

[11] C. Koch, B. Mathur, Neuromorphic Vision Chips, IEEE Spectrum, vol. 33, no. 5, pp. 38-46, 1996.

[12] J. Kramer, R. Sarpeshkar, C. Koch, Analog VLSI motion discontinuity detectors for image segmentation, IEEE International Symposium on Circuits and Systems, vol. 2, pg. 620-623, 1996.

[13] N. Martin, N. Franceschini, Obstacle avoidance and speed control in a mobile vehicle equipped with a compound eye, Intelligent Vehicles '94 Symposium, Proc. of the 1994 IEEE, 1994.

[14] R. Molland, Minimalism in low-level autonomous systems: Bridging the gap between pure reaction and planning, `http://www.psychology.nottingham.ac.uk/research/ai/publications/first-yr-paper-v2/first-yr-paper-v2.html`.

[15] J. A. Moya, G. W. Donohoe, M. J. Wilcox, Cellular modeling of a fly photoreceptor, 1997 IEEE International Conferenceon Systems, Man and Cybernetics, Computational Cybernetics and Simulation, IEEE, 1997.

[16] F. Mura, N. Franceschini, Obstacle avoidance in a terrestrial mobile robot provided with a scanning retina, Intelligent Vehicles Symposium, 1996, Proc. of the 1996 IEEE, 1996.

[17] X. Qi, R. B. Northrop, Dynamic properties of the clock-spike system of the fly, IEEE Engineering in Medicine & Biology Soc. 11th annual Conference, 1989.

[18] S. Viollet, N. Franceschini, Biologically-Inspired Visual Scanning Sensor for Stabilization and Tracking, Proc. of the 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems, 1999.

[19] M. J. Wilcox, D. C. Thelen, A retina with parallel input and pulsed output, extracting high-resolution information, IEEE Transactions on Neural Networks, vol. 10, no.3, May 1999, 1999.

[20] A. Yakovleff, D. Abbott, X. T. Nguyen, K. Eshraghian, Obstacle Avoidance and Motion-Induced Navigation, Proc. of the Computer Architectures for Machine Perception Workshop 1995, pp. 384-393, 1995.

[21] A. Yakovleff, A. Moini, From sensing to perceiving: an overview of the bugeye project, Proc. SPIE, vol.2950, p.86-96, 1996.

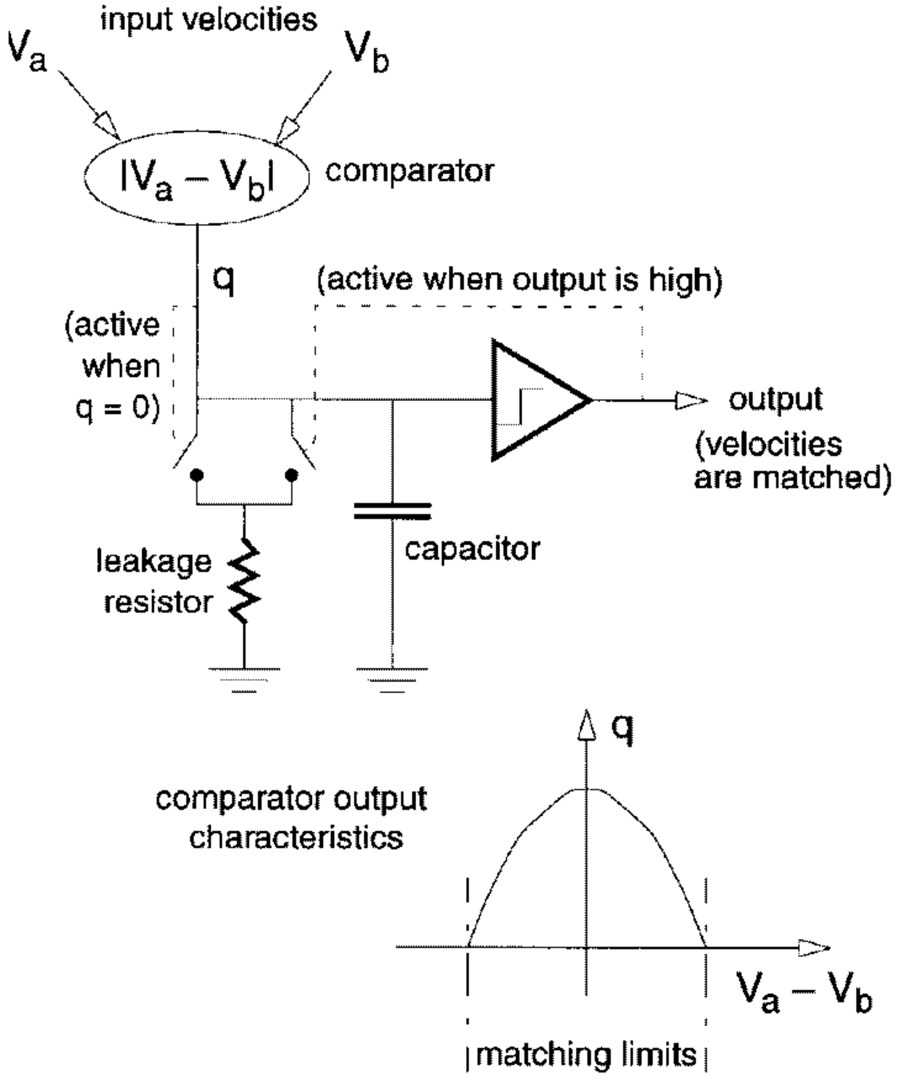

Figure 6: Example of an analog circuit that could be used for velocity matching (see text) [20].

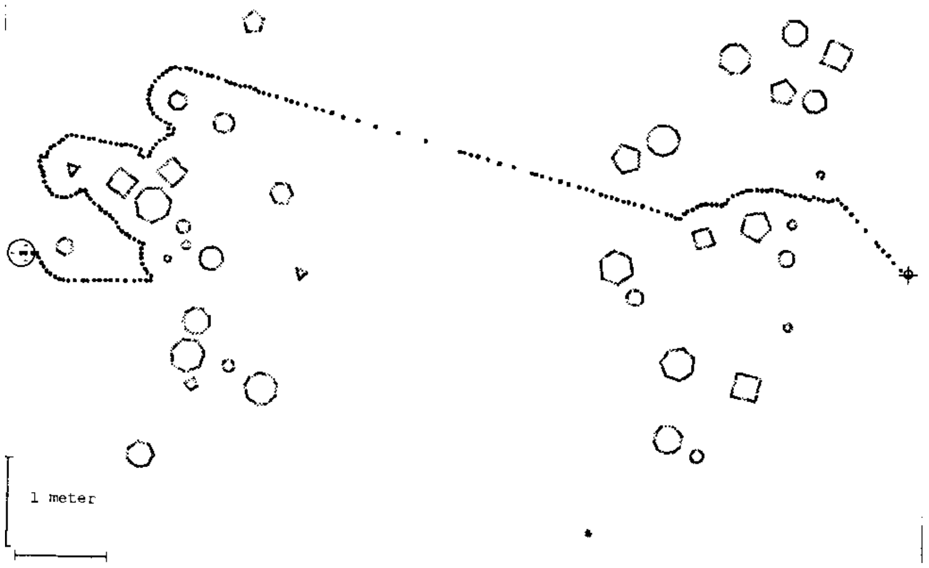

Figure 7: Impressive performance by Franceschini's LMD/EBC implementation (see text) [13].

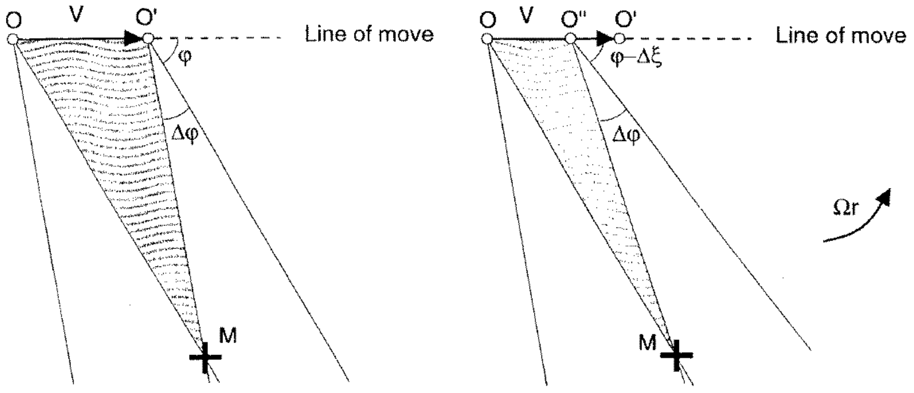

Figure 8: Demonstration of the improvement in angular resolution achievable by using a scanning sensor [16].

Table 1: Useful web addresses for information on fly vision and engineering applications.

Bug-EYE group
`www.eleceng.adelaide.edu.au/Groups/GAAS/Bugeye`

Homepage of M. Franz
`www.franet.de/index.html`

Images of the visual system
`http://brain.biologie.uni-freiburg.de/Atlas/text/VisualFi.html`

Home Page of the Koch Laboratory / CalTech
`www.klab.caltech.edu`

Images, genes etc. of Drosophila
`http://iubio.bio.indiana.edu:81/IUBio-Software+Data/flybase/allied-data/interactive-fly/aimain/1aahome.htm`

Student work on coupling between photoreceptors
`http://www.gel.ulaval.ca/~vision/recherche/etudiants/Gourdeau_C_a.html`

CNRS / Laboratoire de Neurobiologie
`http://irlnb.cnrs-mrs.fr/`

Homepage of Reid Harrison / Ph. D. thesis
`http://www2.elen.utah.edu/~harrison/`

Hans van Hateren's lab / Vision Group
`http://neuro.phys.rug.nl/vision/program.html`

Report on the micromachined gyroscope modeled after the fly's halteres
`http://www.erc.caltech.edu/Research/Reports/landolt2full.html`

Micromechanical Flying Insect (MFI) Project / Berkeley University
`http://robotics.eecs.berkeley.edu/~ronf/mfi.html`

Biomimetics Robotics / Stanford University
`http://www-cdr.stanford.edu/biomimetics/`

The Dickinson's lab
`http://socrates.berkeley.edu/~flymanmd/`

Research report: Visual Sensor with Resolution Enhancement by Mechanical Vibrations
`http://www.erc.caltech.edu/Research/Reports/landolt1full.html`