Bio-inspired motion detection in an FPGA-based smart camera module

T Köhler[1], F Röchter[1], J P Lindemann[2], R Möller[1]

[1] Computer Engineering Group, Faculty of Technology, Bielefeld University,
33501 Bielefeld, Germany
[2] Department of Neurobiology, Faculty of Biology, Bielefeld University,
33501 Bielefeld, Germany

E-mail: `tkoehler@ti.uni-bielefeld.de`

Abstract. Flying insects, despite their relatively coarse vision and tiny nervous system, are capable of carrying out elegant and fast aerial manoeuvres. Studies of the fly visual system have shown that this is accomplished by the integration of signals from a large number of elementary motion detectors (EMDs) in just a few global flow detector cells. We developed an FPGA-based smart camera module with more than 10 000 single EMDs, which is closely modelled after insect motion-detection circuits with respect to overall architecture, resolution, and inter-receptor spacing. Input to the EMD array is provided by a CMOS camera with high frame rate. Designed as an adaptable solution for different engineering applications and as a testbed for biological models, the EMD detector type and parameters like the EMD time constants, the motion-detection directions, and the angle between correlated receptors are reconfigurable online. This allows a flexible and simultaneous detection of complex motion fields like translation, rotation and looming, such that various tasks e.g. obstacle avoidance, height/distance control, or speed regulation can be performed by the same compact device.

1. Introduction

1.1. Models of Insect Motion Detection

Reliable detection of visual motion is needed for several applications in mobile robotics. Especially for obstacle avoidance and flight stabilization it is important to recognize ego-motion. This task is accomplished easily by insects in spite of their small number of neurons and coarse visual resolution. For example, the housefly *Musca domestica* has only about 3000 ommatidia (Land 1997).

Based on behavioural experiments, Hassenstein & Reichardt (1956) suggested the model of an Elementary Motion Detector (EMD), a rather simple signal processing element where the signals of only two photoreceptors are compared to detect a motion in one "preferred" direction. By integrating over many of the EMD responses (with different preferred directions), various types of motion can be detected in a large field of view. In the original EMD model, the signals from the photoreceptors in the ommatidia (i.e. facets of the compound eye) are delayed and correlated with the undelayed signals of nearby receptors, see figure 1(a) (Hassenstein & Reichardt (1956), review: Egelhaaf & Borst (1989)). The delays are carried out by first-order low-pass filters (one per half-detector) with a time constant in the range of 35 ms to 500 ms in the original model (e.g. Harris & O'Carroll 2002). As a correlation function, a simple multiplication yields good results, corresponding well with biological data. In the final summation of the two half-detectors, opposite signs have to be employed, i.e. in figure 1, a motion from the left to the right would lead to a positive response, a motion from the right to the left would result in a negative response in each of the three EMD types. The output signals of multiple EMDs are summed, leading to an averaging over large receptive fields. The preferred directions of the EMDs (resulting from the spatial arrangement of the two correlated receptors) determine to which flow fields the EMD array is optimally sensitive. For example, to detect linear motion in a single direction, all summed EMDs have the same preferred direction. To detect more complex types of visual motion (like rotation or looming), the preferred directions of the summed EMDs form typical patterns as required (see figure 3(a) and figure 9(a)).

As Harris et al. (1999) pointed out, there are some differences between empirical data and the prediction of the standard Hassenstein-Reichardt model. For example, for certain changes in the spatial frequency of presented patterns, the neural output shows a variability that can be explained by the formation of afterimages in the visual pathway (Harris & O'Carroll 2002). To reproduce the observed effects, a first-order high-pass filter is introduced in the model. This extension filters the photo-receptor signals and suppresses constant illumination differences (figure 1(b)). High-pass time constants from 50 ms to 500 ms were used by Borst et al. (2003).

A further variation can be found in the implementation proposed by Kirschfeld (1972) where the high-pass filters are only included in the cross-correlation, i.e. the low-pass filters get the original receptor signals as input (figure 1(c)). The two main effects of this modification are (1) a stronger response and a slightly higher preferred temporal frequency compared with a Harris-O'Carroll detector with the same time constants, and (2) the ability to give an impulse response at the output (Borst et al. 2003). Borst et al. propose another version based on the Kirschfeld detector. In their model, the time constant of the high-pass filters is adaptive and depends on image velocity and image contrast.

Several authors suggest modifications outside the EMD filter core: an improved
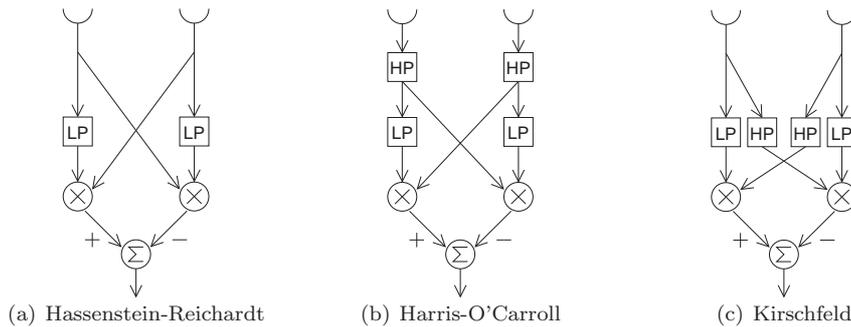
Figure 1. Three versions of EMDs. "∪" denotes a receptor, "LP" a low-pass filter, "HP" a high-pass filter, "×" a multiplication, and "Σ" a summation. Besides the shown model components, several extensions are common. E.g. the modelling of a receptor might be realized by an additional low-pass filter (in a simple model, compare Lindemann et al. 2005) or a series of filters (including light adaptation, see van Hateren & Snippe 2001).

preprocessing with temporal filters (Juusola et al. 1995) or a nonlinear passive membrane model rather than a linear summation of EMD outputs (Borst et al. 1995). Lindemann et al. (2005) compared different models with neuronal responses and found a close correspondence of the more elaborated models with the biological counterpart, especially with respect to the saccadic strategy of flight and gaze control of blowflies.

To summarize, there are several slightly different biological models that fit quite well with the empirical data. When applying the EMD model to engineering tasks, depending on task and environment, some models may produce better results than others. Also, model parameters such as the time constants and the inter-receptor angles may be task-dependent. Therefore, the implementation of multiple of the proposed models and the adjustability of the parameters is desirable for an engineering solution suitable for a large variety of tasks or environments.

1.2. Engineering Implementations

To implement an EMD-based flow detection in a technical system, several technologies can be adopted (this list is exemplary rather than exhaustive):

(1) Implementation in software, either on a PC (Zanker & Zeil 2002) or on a microcontroller (Ruffier et al. 2003)

(2) Implementation in (dedicated) digital, analogue, or hybrid hardware, either based on VLSI or aVLSI ASICs (Harrison 2000, 2004, 2005) or implemented with standard ICs (Franceschini et al. 1992)

(3) Implementation in configurable hardware like FPAA (Field Programmable Analogue Array, Ruffier et al. 2003) or FPGA (Field Programmable Gate Array, Aubepart et al. 2004, Zhang et al. 2008)

The first mentioned possibility has either the drawback of a high power consumption (PC) or does not have a suitable performance (microcontroller). The implementation in an ASIC, in contrast, can provide both a high performance and the best energy efficiency, but the drawback of an ASIC implementation is its reduced flexibility. Consequently, when considering the option of being able to select different

EMD models (or to implement new versions) and to change the time constants and other parameters, the ASIC implementation is not an appropriate solution. The same holds for an implementation with standard ICs. Finally, today FPAAs are not available in a sufficient complexity for larger EMD arrays. So, many FPAA chips would be needed (about two to three per EMD, in case of a current Anadigm AN2xxE04 FPAA) — leading to large circuits and a high power consumption. Thus the implementation based on FPGAs turns out to be the best way for realizing adaptable EMD arrays of a high resolution.

The FPGA-based implementation by Aubepart et al. (2004) achieves very high frame rates of 2.5 kHz up to 5 kHz, but with only 12 photo-diodes as receptors. Their FPGA design is capable of computing 245 EMDs at a frame rate of 2.5 kHz. Zhang et al. (2008) present an FPGA-based EMD array which receives visual input from a CMOS camera with 256×256 pixels. This solution achieves a high frame rate of 350 fps and a very low latency for the EMD computation of 0.25 μs (plus 2 ms for the image transfer). However, a drawback of this implementation might be its size and power consumption; it is implemented on a PCI-FPGA-card in a host PC with an external camera connected via Camera Link.

2. Current Implementation

Our goal was to design a flexible and freely configurable smart camera module that can be used for two different purposes: First, the camera module should serve as a testbed for studies of different biologically inspired motion detection models. Second, the module should be usable as a smart and stand-alone motion detection module for several tasks in mobile robotics. To fulfil these purposes, the smart camera module needs to provide a large number of EMDs, different types of EMDs (as described in section 1.1), inter-receptor angles and time constants comparable to different insect species, and multiple "integrator neurons" to simultaneously detect different motion fields. Furthermore it needs to supply extensive adjustment and control facilities and (for testing purposes) detailed data output such as the current camera image, the current filter results, and the current optical flow field. Finally, future extensions like a second camera or more elaborated EMD versions should be possible. Because of these demands we decided to implement the motion detection module based on a standalone FPGA module with a high-speed CMOS camera and a four-channel DAC for an analogue output of the integrator signals.

The system consists of two main functional units and two peripheral blocks (figure 2(a)). The camera, the DAC, the serial interface circuit, display and configuration switches are mounted on one PCB, and the FPGA together with Flash and DDR-SDRAM memory on a second PCB, stacked as a compact two-board module. Figure 2(b) shows the module which has a compact size of 105 mm $\times$ 57 mm $\times$ 50 mm and a moderate power consumption of about 1 W.

The camera is a Kodak high-speed CMOS camera (KAC-9630) with a resolution of 127×100 active pixels, a maximum frame rate of 580 fps, and an 8 bit parallel interface. Via an I^{2}C interface, the integration time and an analogue gain can be adjusted. When using the camera under dark illumination conditions (e.g. indoors), a long integration time is needed. This leads to a small frame rate (e.g. 100 fps in our laboratory room). Under brighter conditions (e.g. outdoors in daylight), the integration time can be set to shorter values and therefore higher frame rates are possible. With the current FPGA design, the frame rate can be set online to up to
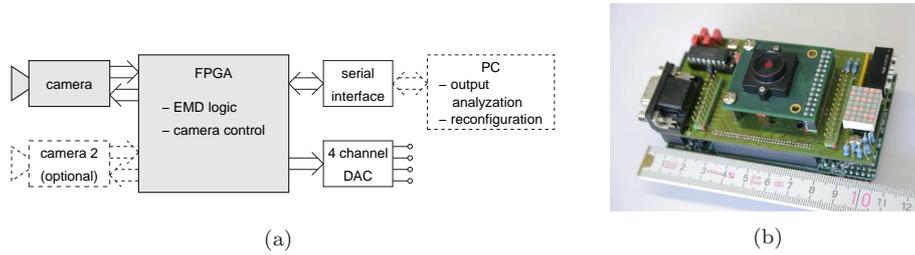
Figure 2. (a) Schematic of the EMD module. The dashed components are not part of the module. Main functional units are shaded. (b) EMD module, camera side (ruler units: cm).
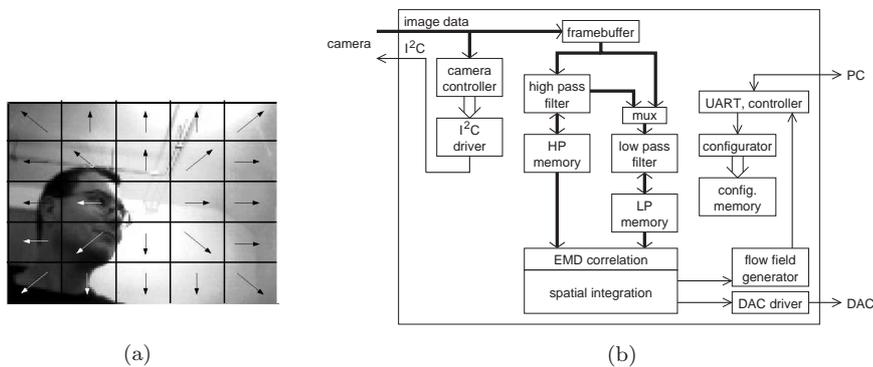
Figure 3. (a) Configuration example: Looming detector. (b) Schematic of the main components inside the FPGA. Thick arrows denote high bandwidth (image) data.

200 fps. Higher frame rates up to 580 fps are in principle possible but need a redesign. This was not yet tested.

Through the integration time, the camera shows properties comparable to a temporal low-pass filter. The effects when changing the integration time can be compared to the effects of changes of a low-pass filter's time constant. Thus, the peripheral filters modelling the biological photoreceptors are implemented in the smart camera module, too. However, a low-pass filter is still a simple model of the biological preprocessing (see section 1.1 and figure 1).

For the computation of the EMDs, the camera images are divided into 5×5 groups of 25×20 pixels, each group with an independently online-selectable preferred detection direction (one out of eight directions in $45°$ steps or "off", see figure 3(a)). For each of the four DAC output channels one such configuration is stored. The pixel distance between the correlated receptors of a single EMD (i.e. the inter-receptor angle) can be configured in the range from 0 to 10 by a parameter called a_{xx}, see figure 4. With the currently mounted lens (horizontal field of view $40°$), the inter-receptor angle is adjustable in the range of $0.31°$ to $3.1°$. Land (1997) compared the interommatidial angles of several insect species and reports that the minimum interommatidial angle
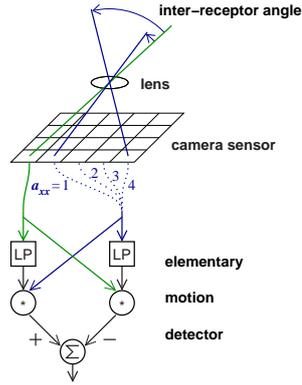
Figure 4. The inter-receptor spacing can be adjusted via a parameter a_{xx} in the range from 0 (EMD inoperable) to 10 pixels. The corresponding inter-receptor angle depends on the installed lens.

varies from tens of degrees down to $0.24°$ and that many common flying insects (e.g. bees and flies) have interommatidial angles in the range $1°$ to $3°$ – which is covered by our camera module. Depending on the configured detection directions and receptor distances, the maximal number of EMDs varies slightly (e.g. 10 000 "left"/"right" EMDs with an inter-receptor spacing a_{xx} of 5 pixels, the maximum is $5 \times 5 \times (20 \times 24) = 12\,000$ EMDs with $a_{xx} = 1$ and a horizontal detection direction).

The EMD logic is implemented completely on the FPGA, an Altera Stratix II S60. The current EMD implementation uses only 13% of the logic elements and 39% of the FPGA internal memory, leaving enough space for further extensions. Figure 3(b) shows the main circuit blocks inside the FPGA. The camera controller analyzes the current image and controls the camera parameters. It can either automatically select the integration time and the CMOS camera analogue gain, or fixed values can be selected manually.

The EMD computation is essentially sequential (with the main data flow from the top left to the bottom right in figure 3(b)), except that the different computational units (high-pass, low-pass, EMD correlation and integration circuits) run as a pipeline in parallel. The high-pass and low-pass filters are first-order types. The two filter circuits run in parallel to each other and compute the filtered images while the pixel data is being read from the camera. A multiplexer at the input of the low-pass filter ("MUX" in figure 3(b)) is used to switch between the Harris-O'Carroll and the Kirschfeld configuration. The Hassenstein-Reichardt detector type is feasible in both configurations by bypassing the high-pass filters.

The "EMD correlation" block consists of four computation circuits that compute four EMDs in parallel for each location in the image (with the preferred directions: $0°$, $45°$, $90°$, $135°$). The four correlation results are accessible for the spatial integration which is done for $4 + (25 \times 4)$ "integrator neurons" in parallel: (1) four configurable flow fields for the four DACs, and (2) $(5 \times 5) \times 4$ (fixed) flow fields for a computation of the current optical flow. The output of the four configurable integrators is amplified by an adjustable gain of 2 to 256 and sent to a low-pass filter (e.g. to simulate the temporal properties of the HSE cell). If this filter is not needed (as in all of our tests) it can be bypassed. For each of the $(5 \times 5) \times 4$ other fixed integrators (with the flow

(a) (b)
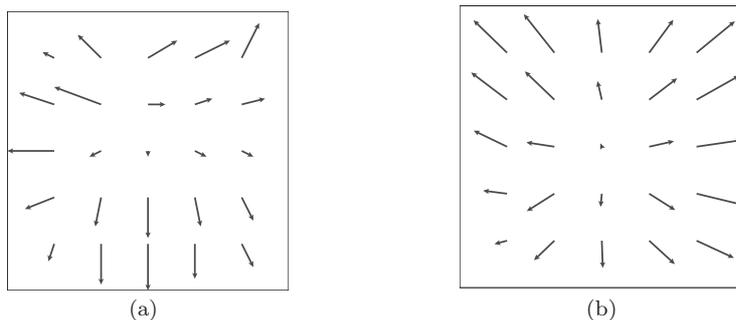
Figure 5. (a) Example of a flow field. Depicted is frame number 135 of 5333 during a series of manual movements of the camera module towards and away from a chequerboard pattern like in figure 10(a). (b) First principal component of the whole set of 5333 frames.

field directions 0°, 45°, 90°, 135°) the responses of all EMDs of a EMD group (in the respective direction) are summed. Thus, the four integrator outputs of one EMD group show the amount of horizontal, vertical, or diagonal motion in the respective EMD group. A vector summation over the four integrator results (in each group) leads to a 5-by-5-sized flow field for the whole image. The current flow field result is accessible via the serial interface.

The EMD computation of the 25 groups and within each group of the EMDs is done sequentially from the top left to the bottom right at a clockrate of 10 MHz. It is started when the first 20×125 of the low- and high-pass results (needed for the first row of five EMD groups) have been updated and continues while the filter circuits are calculating the following image parts.

The module can be used either stand-alone, with limited configuration options via onboard switches and with the four analogue output channels for tapping the four configurable EMD integrator neuron outputs, or in combination with a PC to allow a maximum flexibility in reconfiguration and to store the EMD output or other internal data. Via the serial interface, the EMD configuration (Hassenstein-Reichardt, Harris-O'Carroll, or Kirschfeld type), the time constants of the low-pass and the high-pass filters, the inter-receptor angle, and the peripheral settings (camera integration time, camera gain, EMD output gain) can be adjusted online.

Furthermore, the EMD output, the current camera image, the output of the high-pass or the low-pass filters, and the current flow field (see figure 5(a)) are accessible via the serial interface. A time series of the current flow field can be collected for a given detection task in the respective environment. When analyzing this set of flow fields (e.g. with a principal component analysis) a typical flow field can be computed. This typical flow field, in turn, can be used as a template for the configuration of the EMD array (figure 5(b) shows an example of a looming task).

3. Experiments

Several test setups were used to analyze the module's response to artificial stimuli, three of which will be presented following. In setup I, the EMDs of all 25 regions are configured with the same preferred direction and a linear optical flow in this preferred

(a)   (b)
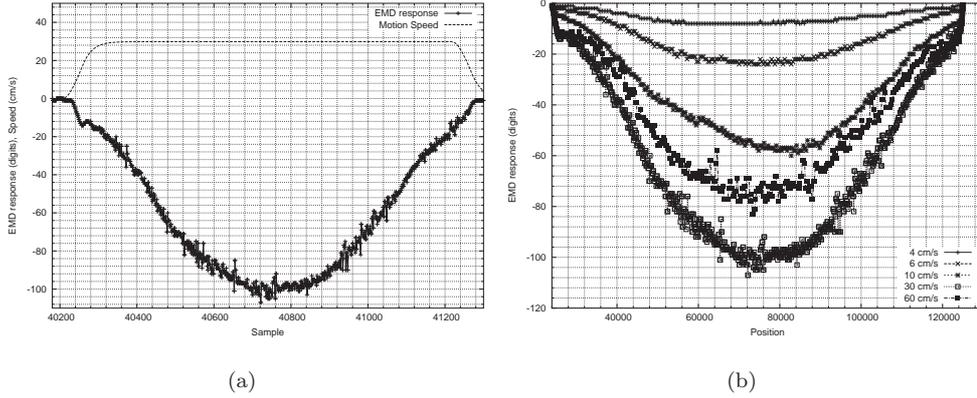
Figure 6. Test setup I. (a) EMD array response, one run. The response is plotted in digits over samples (taken every 5 ms). The module was configured as Harris-O'Carroll detector with the time constants $\tau_{LP} = 35$ ms and $\tau_{HP} = 60$ ms; the configured preferred direction coincided with the negative motion direction. Before and after the printed sine pattern, a white surface is presented to the camera. At the start position and at the end position, only a small part of the sine pattern is visible by the camera. In the centre (about sample 40 700 to sample 40 900), the camera's field of view is completely filled with the sine pattern. The speed in this experiment was 30 cm/s. The motion speed is plotted in cm/s (dashed curve without marks). (b) EMD output of five runs with different speed. The abscissa is the camera position, measured in encoder tics of the robotic gantry's x-axis (1 tic is equivalent to 0.0171 mm).

direction is presented to the camera. Setup II is used to compare the responses of four different detectors while the EMD camera module is being moved in various ways in front of a chequerboard test pattern. In setup III, the smart camera module is tested on a balancing robotic platform when configured as looming detector.

3.1. Test Setup I: Linear movement detection

This test was carried out to systematically analyze the response of the EMD array. The camera module was mounted on a 4D robotic gantry crane (three translatory axes, one rotatory axis). The movement was carried out translatory with constant speed in parallel to a stationary printed sinusoidal bar pattern. The sine stimulus had a period length of 40 mm and was surrounded by white background. The distance between the camera and the printout was about 1.27 m.

Figure 6 shows the time responses for a single run (a) and for five runs with different speed (b). The preferred detection direction was configured against the motion direction, in these runs. Thus, in figure 6(a), the time when the sine pattern fills the camera image can be made out by the linearly increasing (negative) output value, while the disappearance of the pattern leads to a linearly decreasing (negative) output value. Below the EMD response, the motion speed is plotted. As can be seen, the speed is kept constant at 30 cm/s while the camera passes the sine pattern.

In figure 6(b), with increasing speed (6 cm/s to 30 cm/s) the average response first increases, too. When the speed is increased further (60 cm/s), the response gets smaller. To study this behaviour, it is common to plot the average steady-state responses as a function of speed or frequency. This is done in figure 7(a) for test runs
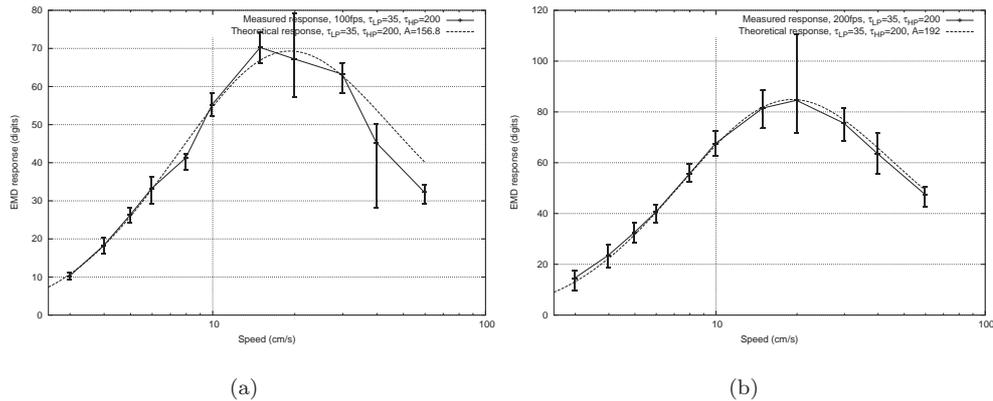
Figure 7. Test setup I. The EMD response is plotted as solid curve over the motion speed (x axes in log scale). The EMD camera module was configured as Harris-O'Carroll detector with a low-pass time constant τ_{LP} of 35 ms and a high-pass time constant τ_{HP} of 200 ms. Additionally, the theoretically expected responses are plotted as dashed curves (as derived by Borst et al. 2003). Since the actual camera input is not available, instead of the theoretical input sine amplitude (ΔI), a free parameter A is used and manually adapted to the measurement data. Therefore, just the curve's shape, peak position, and the relative values can be evaluated, not the absolute ones. (a) frame rate set to 100 fps. (b) frame rate set to 200 fps.

with a Harris-O'Carroll detector, configured with the time constants $\tau_{LP} = 35$ ms and $\tau_{HP} = 200$ ms. Here the typical response behaviour of an EMD can be seen (compare figure 1 in Borst et al. 2003). Borst et al. (2003) derive the theoretically expected steady-state responses for the Harris-O'Carroll detector. As figure 7(a) shows, the measured responses match quite well the expected values. The strongest differences appear at the higher motion velocities. This is supposed to be caused by the temporal aliasing effect (while the implementation is time-discrete, the theory is based on a continuous model). The camera frame rate was set to 100 fps, leading to a Nyquist frequency of 50 Hz. Thus, with the sine pattern's period length of 40 mm, at 200 cm/s an EMD response of 0 should be measured. We therefore repeated the experiment with a higher frame rate, to move the zero crossing to higher speeds. Figure 7(b) shows the responses for a configured frame rate of 200 fps (expected zero crossing at 400 cm/s). Here, even at 60 cm/s the measured average values match the theory.

Further test runs were carried out in another test setup. Here, a sine pattern with a period length of 6.9 mm was used and the frame rate was configured to 100 fps. Here, the temporal aliasing effect should lead to a EMD response of 0 at a speed of about 34 cm/s. For even higher velocities, the responses should be negative, which is confirmed by the measurements (see figure 8).

When comparing the Harris-O'Carroll model and the Kirschfeld model, a much higher output signal is observed when the module is configured as a Kirschfeld detector. The three curves in figure 8 have approximately equal peaks but the output gain of the EMD module in the Harris-O'Carroll configuration was set to the double of the amplification in the Kirschfeld configuration ($16\times$ instead of $8\times$).

The influence of the time constants can be seen when comparing the plots for the two Harris-O'Carroll configurations in figure 8. As Borst et al. (2003) states, when
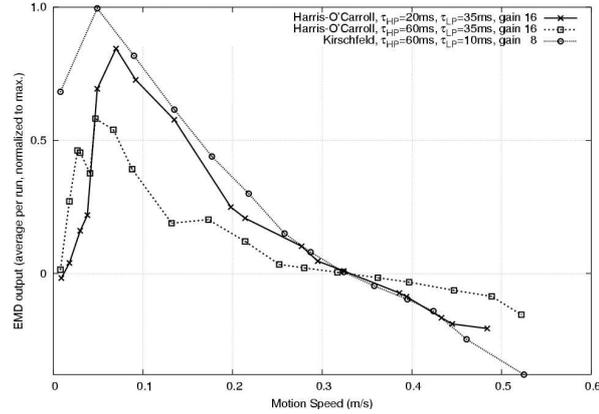
Figure 8. Average EMD output, test setup I. τ_{HP} and τ_{LP} denote the time constants of the high-pass and the low-pass filters. These test runs were carried out in a modified version of test setup I (not with the robotic gantry crane). The period length of the sine pattern was 6.9 mm. The distance between the camera and the sine pattern printout was about 20 cm. As EMD array, 500 EMDs were configured.

keeping the low-pass time constant unchanged, then the larger the high-pass time constant, the more the amplitude spectrum shifts towards slower velocities.

3.2. Test Setup II: Comparison of Different Detectors

In the second test setup, the responses of four different detector field configurations are compared. The first integrator was configured as linear vertical detector, the second as linear horizontal detector, the third as looming, and the fourth as rotation detector (see figure 9(a)). The camera was exposed to a printed chequerboard test pattern (figure 10(a)), and undergoing different types of movements (carried out by the four axes of the robotic gantry used in test setup I). In figure 9(b), the responses are depicted.

First, the camera was moved downwards (samples 50...700) then upwards (samples 750...1400). Second, a translatory movement to the right (samples 1400...2200) and to the left was carried out (samples 2250...3050). Third, the camera was moved towards (samples 3100...3860) and away from the pattern (samples 3920...4700) and finally a counterclockwise (samples 4750...5170) and a clockwise rotation was performed (samples 5220...5650).

In the first plot in figure 9(b), a strong response of the vertical motion detector to the vertical movement can be recognized. The response level depends on the number of edges visible during the motion (see figure 10(a)). Also the response of the horizontal detector (second plot) shows a noticeable peak during the horizontal motion phase. The amplitude of the looming detector output signal (third plot) is smaller than the signal amplitudes of the linear motion detectors (as the centre EMD group is not used and as the local visual motion velocities differ). However, the detector's response is observably stronger during the movement towards and away from the pattern than during other or no movements of the camera. The rotation detector (fourth plot) generates a clear response to the rotational movement. A rather strong noise can be

(a)   (b)
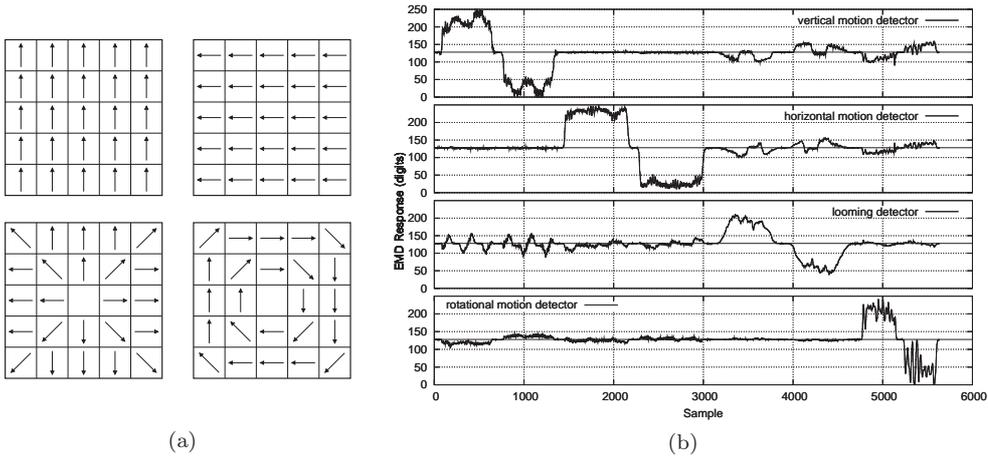
Figure 9. Test Setup II. (a) The configured preferred directions of the four detectors used in test setup II (vertical, horizontal, looming, and rotational). (b) EMD output of 4 simultaneously computed detectors.

(a)   (b)
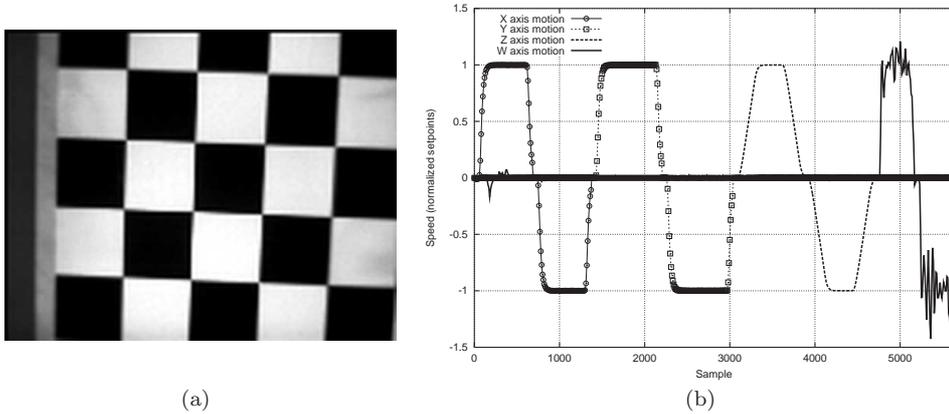
Figure 10. Test Setup II. (a) Image of the printed chequerboard test pattern taken by the camera of the EMD module. Integration time and analogue gain of the camera were chosen to get a high contrast. (b) The motion speed of the four gantry axes in test setup II. The data was generated during the same run as used for the data plotted in figure 9(b). The x-, y-, and z-axis of the gantry is parallel to the camera image's vertical, horizontal, and optical axis, respectively. The w-axis of the gantry is parallel to the camera's optical axis.

seen here, especially during the final clockwise movement. However, in contrast to the other axes, the rotational axis motion is affected by the cable transport. The EMD array response follows nearly perfectly the variations of the motion speed of the rotation (compare figure 10(b)).

The small responses of the horizontal and the vertical detector during the looming phase, and the small responses of the looming detector in the first motion phase are caused by an asymmetry in the image of the chequerboard pattern (see figure 10(a): The number of edges in one half of the image may be different from the number of

Figure 11. Photograph of the two-wheeled balancing robot platform used as test setup III.

edges in the other half).

3.3. Test Setup III: Looming Detection On A Robotic Platform

In the third test setup, the module's performance on a two-wheeled, balancing robotic platform is tested (see figure 11). Compared to a standard three- or four-wheeled mobile robot, this platform has the special characteristic that in addition to the intended movements of the robot (straight forward or backward, turning left or right), an up-and-down motion caused by the balancing process is present all the time. When applying a motion detection system to such a "difficult" platform the question comes up if the primary function of the system (e.g. obstacle detection) still works. In contrast to the previous setups, this test was not meant for a systematical analysis of the smart camera's properties. In fact, it should show its appropriateness for noisy, real-world environments.

The robot has a height of about 37 cm and the camera was mounted on top of it in a total height of 48.5 cm. Based on the measurements of a laser distance sensor, a state-space controller tries to keep the robot at a given angle (upright for zero velocity, increasing inclination with increasing speed). The robot tilts a few degree forward and backward while driving or standing.

The EMD array of one detector output was configured as looming detector (see figure 3(a)) and the robot was driven towards and away from a printed chequerboard pattern (sized 82 cm by 82 cm) twice. Figure 12(a) shows an image taken by the camera module before a run; in figure 12(b), the current robot speed and the respective EMD output can be seen.

In the lower plot, the different phases of the test are distinguishable: Up to about sample 500 the robot goes forward (phase 1), then backward up to sample 900 (phase 2), and then again forward with a little higher speed (phase 3), stopping at about sample 1150 and finally going backwards and stopping slowly (phase 4).

In the EMD response of the first approach, an increasing peak amplitude occurs, up to about 100 samples before stopping (i.e. approximately up to sample 450). In

(a)   (b)
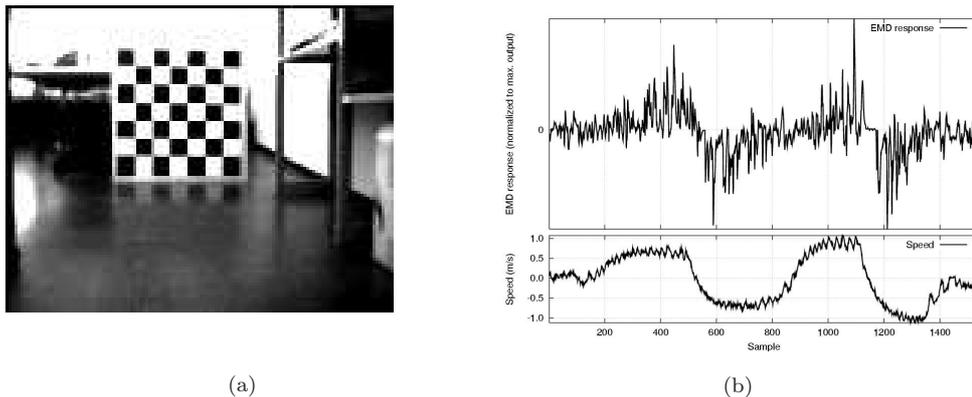
Figure 12. (a) Image taken by the smart camera module during a run in test setup III. (b) EMD output and robot speed, test setup III. The detector was configured as looming detector, the EMDs in the Kirschfeld configuration. The samples were taken every 5 ms.

the second run (in phase 3), the amplitude increases up to about 50 samples before stopping. In both approach phases 1 and 3, after one major peak, the peak amplitudes decrease again relatively fast. In the second run (at the end of phase 3) the EMD response decreases down to a zero output signal when the robot is closest to the pattern (at about sample 1150). When going backwards (phases 2 and 4) the detector response is nearly mirror-inverted. Here, at first the peak amplitudes increase relatively fast (the response is now (mostly) negative); after one major peak, the peak amplitude is slowly decreasing again.

The balancing tilt movements of the platform can be identified in the oscillations of the speed curve. Also the relatively strong noise (when comparing for example with figure 9(b)) is probably caused (at least partly) by this platform's special characteristic. These measurements in general are comparable with results in previous publications, for example the aVLSI collision detector by Harrison (2004), especially also the negative outputs within the phases 1 and 3.

4. Discussion

First, the experimental results show that the motion-detection module described in section 2 generates output signals as one would expect from the respective EMD model and as are described in former publications. Especially the experiments in setup I produced curves that match well with theory (see figure 7) and specifically show the expected output at higher frequencies, resulting from the aliasing effect (see figure 8).

Second, especially the test setups II and III show that the smart camera module can be used for different tasks (and also different tasks in parallel) and even on a difficult, noisy setup like a balancing robot (setup III). However, all these tests use (at least partly) artificial stimuli. Testing in natural environments (as well as the proper adaptation to it) needs some further work. Since the smart camera module allows to use different biologically inspired models and different settings, we expect that an adaptation to other environments is possible. Especially an extension with additional

peripheral filters (linear or nonlinear) is possible and allows a closer adaptation to the biological preprocessing stages.

Third, these experiments lead to some conclusions regarding the models themselves and their application. The Kirschfeld detector generates higher output amplitudes than the Harris-O'Carroll configuration, so it might be the first choice for some practical applications (e.g. obstacle avoidance, see figure 12(b)). On the other hand, its output seems to be more sensitive to noise in the input signal (as figure 12(b) indicates). For some applications, this high sensitivity may be demanded, for others it may not be acceptable. A second potential drawback was visible in other test runs (not included here): With its strong response behaviour and the gain chosen appropriately to distinguish minor peaks, the Kirschfeld configuration generated larger peaks that resulted in a ceiling effect (saturation) at the DAC. Here, more elaborated EMD configurations with a nonlinear spatial integration could improve the detector performance.

5. Conclusion

A flexible motion-detection camera with a high frame rate as well as a high resolution was presented. The first experiments have shown that it can serve as a testbed for different motion detector types with varying configuration parameters and test conditions. Based on the successful results of the presented first tests we are currently carrying out further extensive tests on the Kirschfeld and the Harris-O'Carroll detector to analyze the influence and the optimal choice of the different parameters in combination with different robotic platforms. We will also test further EMD configurations (as described at the end of section 1.1) with the robotic gantry setup to compare the different architectural characteristics when stimulated with artificial and natural test patterns. Here, different extensions in the peripheral preprocessing and especially nonlinear extensions in the spatial integration of the single EMD responses promise improvements for the different robotic platforms and for different engineering applications. With the findings of these systematic tests with artificial stimuli we can directly switch to the mobile robot setup and use the same device in a natural indoor or outdoor environment to solve engineering tasks, benefiting from biological insights.

References

Aubepart F, Farji M E & Franceschini N 2004 *Industrial Electronics, 2004 IEEE International Symposium on* **1**, 71–76.
Borst A, Egelhaaf M & Haag J 1995 *Journal of Computational Neuroscience* **2**, 5–18.
Borst A, Reisenman C & Haag J 2003 *Vision Research* **43(11)**, 1309–1322.
Egelhaaf M & Borst A 1989 *Journal of Optic Society of America A* **6(1)**, 116–126.
Franceschini N, Pichon J & Blanes C 1992 *Philosophical Transactions of the Royal Society of London Series B - Biological Sciences* **337(1281)**, 283–294.
Harris R A & O'Carroll D C 2002 *Vision Research* **42(14)**, 1701–1714.
Harris R A, O'Carroll D C & Laughlin S B 1999 *Vision Research* **39(16)**, 2603–2613.
Harrison R R 2000 An analog VLSI motion sensor based on the fly visual system PhD thesis California Institute of Technology Pasadena, California. http://www.ece.utah.edu/˜harrison/thesis.pdf.
Harrison R R 2004 *in* S Thrun, L Saul & B Schölkopf, eds, 'Advances in Neural Information Processing Systems 16' MIT Press Cambridge, MA pp. 987–994.
Harrison R R 2005 *IEEE Transactions on Circuits and Systems* **52(11)**, 2308–2318.
Hassenstein B & Reichardt W 1956 *Zeitschrift für Naturforschung* **11(9-10)**, 513–524.
Juusola M, Weckström M, Uusitalo R O, Korenberg M J & French A S 1995 *Journal of Neurophysiology* **74**, 2538–2547.

Kirschfeld K 1972 *in* R Wehner, ed., 'Information processing in the visual systems of arthropods' Springer Berlin, Heidelberg, New York pp. 61–74.

Land M F 1997 *Annual Review of Entomology* **42**, 147–177.

Lindemann J P, Kern R, van Hateren H H, Ritter H & Egelhaaf M 2005 *The Journal of Neuroscience* **25(27)**, 6435–6448.

Ruffier F, Viollet S, Amic S & Franceschini N 2003 *in* 'Proceedings of the 2003 International Symposium on Circuits and Systems, 2003. ISCAS '03.' Vol. 3(III) pp. 846–849.

van Hateren J H & Snippe H P 2001 *Vision Research* **41**(14), 1851–1865.

Zanker J M & Zeil J 2002 *in* H Bülthoff, S.-W Lee, T Poggio & C Wallraven, eds, 'Proceedings of the Second International Workshop on Biologically Motivated Computer Vision (BMCV) 2002, Tübingen, Germany, November 22-24, 2002' Vol. 2525 of *Lecture Notes in Computer Science* Springer-Verlag pp. 146–156.

Zhang T, Wu H, Borst A, Kuehnlenz K & Buss M 2008 *in* 'IEEE International Conference on Robotics and Automation (ICRA)' Pasadena, CA pp. 335–340.