Integration of Visual Processes for Control of Fixation

Claus Siggaard Andersen* **and Henrik Iskov Christensen****

Laboratory of Image Analysis, Aalborg University,
Fr. Bajersvej 7D
DK-9220 Aalborg East
Denmark

Abstract. This paper presents an architecture for controlling an agile camera head with multiple degrees of freedom. The purpose of the architecture is to provide control of the different degrees of freedom in the head by applying it in a fixation based scheme. This includes control of the optical degrees of freedom, an often neglected property in most current head control architectures.
The basic thesis is that a tight connection between the lowest visual processes and the sensing apparatus is imperative for successful operation in dynamic environments. This has lead to the introduction of a basic system which facilitates three low level processes, which are: fixation, tracking and attention selection and shifting, which can provide a fully data driven response. The basic system carries the possibility of servicing requests from higher level visual processes, whereby model driven responses in a hypothesis verification/rejection scheme may be implemented. The interaction between the basic system and higher level processes is solely accomplished through the attention mechanism, whereby arbitration between multiple high level processes is resolved.
Using inspiration from the motoric system in the primate vision system an architecture based on Cyclopean representation of visual information has been designed. This is an equal eye dominance approach, which separates control of the fixation point into control of *vergence*, *version* and *tilt* angles, whereby the control of range and "position" is separated, enabling a very simple integration of multiple depth cues for control of the vergence angle.
Based on the experience obtained with the designed and implemented system it is asserted that the proposed architecture provides an efficient way of organizing the lower levels of a fixation based vision system.

1 Introduction

Many of the results that have been achieved through the study of *biological* vision have been used as a basis for *computer* vision. This is particularly true for the studies of primates visual system, with research in eye movements as one of the main contributions to the field. The knowledge obtained from biological systems has lead to research in the use of active perception and *active (computer) vision*, where the sensing apparatus is involved actively in the perceptual process. These ideas have later been formalized into what has become known as *active vision*, by Aloimonos et al. [1] and others.

Active vision has existed as a theoretical discipline for more than 10 years. It was not, however, until the mid-eighties that the required hardware became available, enabling a more empirically founded approach to the field. Among the very first to experiment with active vision ideas were Krotkov [2] and Ballard and Brown [3, 4, 5, 6]. They built vision systems based on the ideas of active vision in combination with a controllable sensor. Using these new systems they proved that real time tracking as well as 3D structure and motion estimation in dynamic scenes was feasible.

Many others have since then designed and constructed controllable sensory platforms, often termed camera heads, or just heads, due to their resemblance with the primate visual system, which often has been the main inspiration for the designers. See e.g. [7] for a survey of current camera head designs.

* Funded by the Danish Technical Research Council under the MOBS framework programme
** Funded by ESPRIT Basic Research Action BR 3038/7108
Email: andersen@cosmos.imag.fr, hic@vision.auc.dk

1.1 Motivation

The diversity between the camera head designs is highly prominent. These differences have mainly been caused by the available hardware at the time of construction, as well as the specific application the camera heads initially were intended for. One of the results of not having a unified camera head design is that the algorithms and architectures for the visual systems differs considerably. That in itself is not a problem, but it hinders the development of a standard approach for system control, which will be more general and easier to extend than many of todays systems. We believe there is a genuine need for a more general framework for designing vision systems - particularly for the lowest levels of an active vision system. Such a framework will enable comparison and transfer of results obtained within the field using different experimental platforms, but should also make it easier to expand initial algorithms by sorting out the hardware control conflicts between individual processes. Thus one may concentrate on the algorithmic development and extensions, rather than being concerned with control which should be facilitated through the architecture itself.

2 The Proposed Approach

Based on the motivation presented in the previous section this project has been concerned with designing a modular architecture for controlling the motoric system of a fixation based camera head with multiple degrees of freedom. This includes control of the optical degrees of freedom, an often neglected property in most current head control architectures.

The fundamental idea is that a tight coupling between the lowest visual processes, referred to as the basic system, and the sensing apparatus, with known latencies, is imperative for successful operation in dynamic environments. But what is the basic system, or equivalently "what is the basic functionality of a camera head"? Considering what has been learned from biological vision, particularly from the primate visual system with its non-uniform retinal layout, processes like *fixation*, *gaze shift*, and *smooth pursuit* seems plausible candidates for a basic foveated system [8, 9]. A system capable of addressing these aspects of active vision will be capable of fixating on an object, and maintaining fixation while it is moving, or during ego motion of the head.

This basic system has the advantage of not requiring a large system for its operation and the data driven response of the system becomes apparent through the design of it. However, this basic system will lack the functionality of performing selection of a fixation point, i.e., the system cannot be initiated, since it has no way of determining what is "interesting". To make the system able to operate without explicitly coding where the target is, and what it looks like, an overt attention mechanism must be added. An overview of the processing modules that are contained in the basic system is shown in figure 1. Note that there is no block in the figure for gaze shift, since this is merely an effect of changing interest (attention).

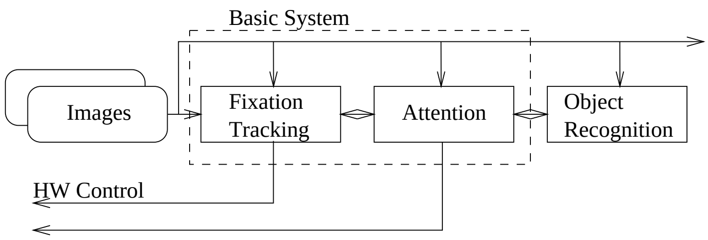

Fig. 1.: The basic system layout with a fixation - tracking module. This is followed by an overt attention module, enabling selection of interesting points or regions in the images. For the complete system an object recognizer is needed for supplying top-down information.

The attentional mechanism will allow for selection of interesting points from the input data. The system can then perform selection of fixation points, fixation and tracking. This basic system does, however, not fully operate in a hypothesis verification/rejection scheme. For this a higher level module is needed, which supplies the top-down information. The claim here is that the higher level module only needs to provide information to the system via the attention module. This has the advantage of providing a simple integration into for example the VAP [10] architecture for a complete system design. The expanded system can thus be controlled using both data and model driven information. It should noted that the system can be used with a rich variety of high-level modules; the object recognition module is merely an example.

So far we have considered which processes to incorporate in the basic system but not how these are to be integrated, though ideas can be derived from related work. During many of the studies of the motoric system in biological vision system control theory has been applied as a tool for analyzing specific motor - sensor interactions. See e.g. [11, 9, 12]. Control theoretic techniques have also been applied when designing computer vision systems, particularly for the lowest levels of visual processing. This is only natural since it here is a matter of controlling a mechanical system such that viewing conditions are optimized for solving a specific task.

In seems thus like a natural choice to use control principles for the design of a low level vision architecture, whereby the architecture can be viewed as a control system linking the motor actions with computational modules. This has two advantages: 1: there are well established techniques for predicting/evaluating the behavior of the system and 2: the simplicity is kept at a level where both design and construction is conceivable.

But what are the building blocks of such a control system, is it a fixation - tracking - attention selection system ?, is it a smooth pursuit - saccadic motion system? or something else?

Any kind of separation criteria could be pursued. However, here the idea is to create a versatile architecture which allows for a separation between control issues and vision algorithms, such that different algorithms may be tested without changing the entire system. Thus a separation according to different visual processes seems natural. On the other hand, three modules doing fixation, tracking and attention selection, do not provide a scheme for interconnection and interface to the hardware, i.e., there is a gap between individual system modules and hardware. In other words, it is necessary to establish an infrastructure between the physical system and the visual processes. From traditional control systems which consist of processing blocks and "control paths", ideas on how to structure the system may be obtained.

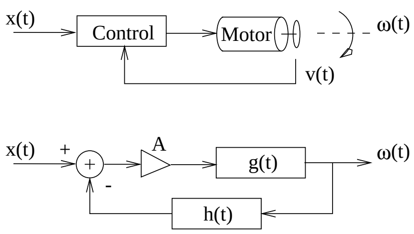

Fig. 2.: Example of standard control system for a DC-motor with tachometer feedback, with "normal" appearance at the top and the control schematic equivalent at the bottom. The driving signal from the controller to the motor is a voltage, proportional to the desired angular velocity. The tachometer readout is a voltage equivalent to the current angular velocity.

In a control diagram as shown in figure 2 all processing and hardware is represented by processing modules. In vision system context this would imply that all visual processes and the sensing apparatus would occupy individual modules. The control paths would then link the individual modules together.

For a traditional control diagram as shown in 2, the paths carry all the information required for controlling the actuator. In this simple case there is only a single connection between neighboring modules, but more complex systems may well have numerous control paths connecting modules.

While the control diagram does not necessarily display the full relation between modules, it does provide an overview, and the general flow of information is visible. Using this representation for a vision system would then mean that the control paths should carry the primary control signals in the system.

3 Designing the Architecture

The "standard" approach to designing architectures for low level active vision systems is to make single visual processes control separate (single) mechanical degrees of freedom. An example of such a procedure was used for the original Rochester head design, which embed real time tracking that relies on separate fixation and tracking. This is achieved by having a leading camera tracking the target and controlling the slaved camera using zero disparity information. Thus tracking and fixation processes are operating simultaneously on separate motors, i.e., left and right pan motor. This is the most common approach reported, which probably is due to its simplicity and generality. Even the Harvard head with mechanically linked cameras such that vergence always is symmetric has employed this technique [13]. The control for doing fixation and tracking is thus fully decoupled, as demonstrated in [6]. More advanced systems embedding attentive mechanisms as well needs an additional control loop. For the Harvard design this was performed by adding an extra layer in the control system performing attentive processing and selection. This higher level module thus produced a saccadic movement to a new area of interest by adding an offset to the current camera motor angles, while halting the lower level visual processes for tracking and fixation.

Vision in primates has, however, been shown to function in a different manner [8, 11, 12]. Yarbus, Carpenter and many others argues that very convincing experiments have proved that eye movements typically are performed in two separate stages, *version* and *vergence*, with both eyes participating in both motion patterns, while fixating a some point in space. The *version* angle is the direction of gaze for an imaginary eye positioned between the two rotation centers in figure 3. The *vergence* angle is the angle between the optical axis of the two cameras as depicted in figure 3. The *tilt* angle denotes the angular elevation of the fixation point from the horizontal plane. This approach is called Cyclopean control, where no single eye has full dominance as opposed to the previous leading control strategy. Also in a computer vision context the equal eye dominance approach has been shown capable of handling monocular occlusion, and facilitating more simple fusion of similar but separate cues. In [14, 15] it was shown that accommodation and disparity cues easily can be fused to perform cooperating control of the two corresponding actuators, vergence angle and lens accommodation distance. It seems therefore that the equal eye dominance approach is an interesting alternative to the standard control strategy, thus it has been chosen as the basis for the proposed architecture.

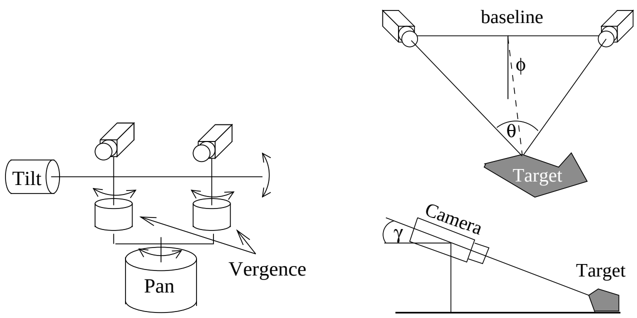

Fig. 3.: Schematic layout of the Aalborg head visualizing the mechanical degrees of freedom. To the left is the layout of the head and to the right the Cyclopean representation of the head with the version ϕ, vergence θ, and tilt γ, angles shown. The version angle relies on the two vergence motor settings. The pan motor contributes however, along with the vergence motors to the direction of gaze. The vergence angle does as the version angle only rely on the vergence motors. The tilt angle depends on a separate motor.

When designing an architecture based on the Cyclopean representation for equal eye dominance control, control of fixation is converted into control of *vergence*, *version* and *tilt* angles (see figure 4). The Cyclopean angles are unfortunately not directly related to individual mechanical actuators. They depend on the combined setting of the two vergence motors, as shown in figure 3. Thus the control system seems to be more complicated than the leading eye counterpart. There exist, however, a simple solution for solving the interdependence [16], and thus this is not really an argument for not choosing a Cyclopean eye representation.

Note that the Aalborg head has a common tilt axis, thus this is a single degree of freedom which may be controlled as in the leading eye approach. There do however exist head designs with separate tilt axes for the two cameras. To utilize this extra degree of freedom horizontal alignment of the the two cameras needs to be considered. This may again be performed by driving one of the cameras as leading and the other following. Typically it is desirable to have zero disparity vertically, thus often such systems only use the extra degree of freedom for an initial alignment between the two eyes.

When considering the Cyclopean representation in the light of the processes chosen for the basic system some conflict seem to arise. While the processes maintains the separate control with a traditional leading eye control strategy, the visual processes, "fixation, tracking and attention selection", does not seem to fit into a cyclopean control scheme. However, the visual process of tracking in the leading eye approach is roughly equivalent to performing control of version and tilt in the cyclopean representation, while fixation corresponds to the process of vergence control. Hence renaming the modules, and utilizing a different representation the basic control architecture may facilitate equal eye control, as shown in figure 4. In this new "domain" being *vergence*, *version*, *tilt*, and *attention selection*, the lower modules are again directly linked to separate degrees of freedom, though these are "pseudo mechanical degrees of freedom". Utilizing the same strategy as in the system proposed by Clark and Ferrier for the Harvard head, the attention process can control the sensing system through these lower level system modules, by simply adding an offset to the current position, whereby abrupt changes may be introduced, see figure 4.

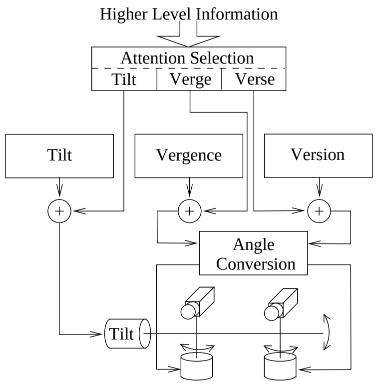

Fig. 4.: Overview of the proposed architecture for an equal eye dominance control scheme. Only the control of angular degrees of freedom have been shown.

It should be noted that the figure only displays the "forward control lines", which has been chosen for readability considerations. Actually there are feedback signals from the hardware to the visual processes, as well as communication between the individual processing modules. The signals in the system is as described earlier the actions issued by the processing modules, which in this case is vergence, version and tilt angle adjustments. Thus the close connection with the actual control of hardware is still maintained, except that the version and vergence angles do not directly map to mechanical actuators. Beside the actual control paths there are of course image path-ways to each module.

3.1 Completing the Architecture

In the previous section we presented a system for controlling the mechanical degrees of freedom, associated with "eye movements". The Aalborg camera head is though in possession of one additional rotational degree of freedom - the pan motor, as shown in figure 3. Beside the pan motor the head is also equipped with motorized lenses, an option most of today heads possess. The lenses do in this case have three degrees of freedom, being focal length (zoom), focus (accommodation) distance, and aperture.

None of the optical degrees of freedom are represented in the control architecture presented previously, which immediately seems like a flaw. This is actually a missing part in most camera head control architectures. The main reason being that the control system is intended for a fixation based system, thus the angular degrees of freedom are the most vital as they are the ones defining, where the cameras are fixated. This means the vergence, version and tilt angles can be viewed as the *primary* degrees of freedom, whereas the optical ones, and the pan angle are of a secondary nature, which may be set to aide the functionality of the primary ones.

The support provided by secondary degrees of freedom for primary ones, is still an unexplored field of research, though a few obvious ones have been considered through the past. The foremost known is fusion of accommodation and disparity depth cues for controlling the vergence angle and accommodation distance. The system architecture presented in the previous section and displayed in figure 4 is not directly able to facilitate this fusion, since it merely considers rotational degrees of freedom. Thus the optical degrees needs to be incorporated as well. The use of these do unlike the primary ones rely extensively on the actual algorithms implemented. For example, the accommodation cue extraction may be based on control of either aperture size or accommodation distance. Furthermore the aperture could also be used as “fidelity” measure, i.e., it could be controlled with the sole purpose of optimizing the viewing conditions.

The choice of the specific use is also extremely dependent on the sensing apparatus. For the Aalborg head the aperture has no (non-visual) feedback and the feedforward control is highly uncertain, thus applying this actuator in the focus ranging is not feasible. Furthermore the focus ranging has as a goal to control the accommodation distance, thus it seems also more natural to apply this actuator in the sensing process. For this reason the aperture seem much more useful for “fidelity” control, where the intention is to ensure optimal viewing conditions. This may be performed by using the video signal to determine the aperture setting in closed loop. This process is known from many camera systems as auto iris, or automatic lighting control, ALC.

It should be emphasized that both types of compensation do change the image intensity locally, thus often auto focusing (accommodation distance measures) is confused. Compensation during accommodation estimation is therefore not desirable.

The last optical degree of freedom, zooming, has not been considered previously in this work, but research using it for extraction of depth information [17, 18] has been reported elsewhere. The actuator also seems useful for attention selection schemes with multi-resolution processing, or simply for getting overview images of the scene, while at the same time having the option of zooming in on interesting locations. The use of zoom in combination with attention selection is believed to be an interesting approach which may provide the spatially non-uniform sampling others are performing in software.

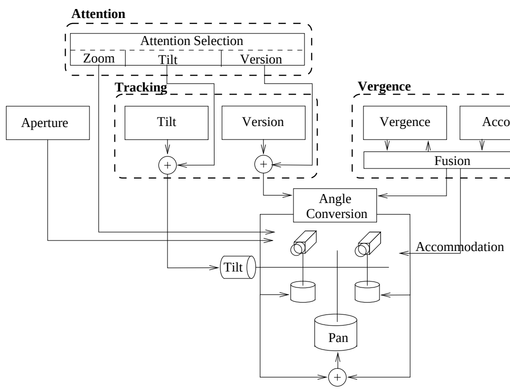

Fig. 5.: The modified cyclopean control architecture where the optical degrees of freedom and the pan motor have been added.

While all the degrees of freedom considered so far have a direct relation with one or more visual processes, the pan motor has not. It may be controlled in numerous way, but it has been decided to control this degree for the sole purpose of minimizing the difference between the left and right vergence motor angles. This has one very important implication, which is that the version angle is minimized, i.e., symmetrical vergence is achieved. The control does not need any visual feedback, since it is only intended for reducing the difference in vergence motor angles. This implies that the pan motor control may be implemented within the hardware controller, as shown in figure 5. The servo signal for the pan motor control is simply the version angle. Whenever it is non-zero, the pan motor compensates by a counter rotation.

The dynamics of the pan motor has to be considered in comparison with the vergence motors to ensure stability. For the Aalborg head this is not a problem, since the pan motor is 10 times slower than the vergence system and has a much higher inertia. This means that a factor of 10 between the compensation and the vergence/version control is inherent in the mechanical system.

4 An Experimental System

Based on the architecture outlined above a complete basic vision system was designed and implemented to get some experience with the architecture. This includes experimenting with different algorithms for controlling the primary and secondary degrees of freedom. Based on the experience obtained through this work, it was chosen to build a final system relying on correlation based image stabilization for the left and right camera. The computed image slip from the two cameras is combined to form the error signal for control of the the version and tilt angles. While a disparity estimate could be computed from the target location in the image pair, it has been chosen to perform an explicit disparity extraction by correlating the images. This provides of course redundant information, but it also allows for a more robust control since a loss of disparity information does not necessarily mean that version and tilt control cannot be performed and vice versa.

The disparity estimate is though not the sole measure to control the vergence angle, as it was stated above. The fusion of multiple cues has been considered important and is particularly simple due to the representation chosen here. For example a fusion of multiple depth extraction cues may be used to control the vergence angle. This has been exploited here, as shown in figure 5, by implementing an integration scheme proposed by Carpenter [12]. This architecture was originally introduced to explain the fusion process between disparity and accommodation cues undertaken in the human visual system for combined control of the lens' accommodation strength and the vergence angle. The principle is that both cues influence the setting of both actuators, though often with different strength, which handles tricky situations like monocular occlusion. See [16, 15] for an in depth explanation.

Using the integration strategy proposed by Carpenter the architecture designed here is capable of handling monocular occlusions, i.e., vergence angle control can be performed even though no disparity information is available.

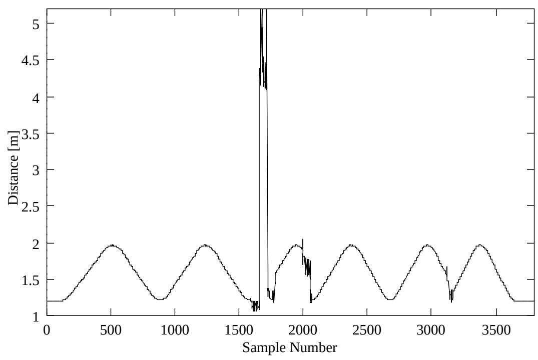

Fig. 6.: Fixation distance for the combined disparity and accommodation control. Normal operation is based solely on disparity, but when no disparity match can be obtained the accommodation process is activated. The objects moves towards and away from the object between 1.2 and $2.1m$.

So far only the "active" components controlling the actuators have been considered. The system also includes an attention selection mechanism, which controls the fixation point by posing step changes in the current Cyclopean angles. The selection process is based on a method proposed by Culhane and Tsotsos [19], but modified to handle multiple feature input and top-down information [16]. Figure 7 shows an example of the entire systems operation on a static scene, where interesting regions (white rectangles) have been selected using the attention mechanism.

The only cue used for the displayed example was edge information, thus the receptive field should be centered around an area with high edge contrast. This is also the case as seen from figure 7, where the 5 most prominent receptive fields have been displayed on top of the image. Note that the long vertical bar in the left side of the image is an effect of image wrapping occurring in the framegrabber, and should not be considered a valid region.

Using the information from the derived receptive fields the attention can be shifted between the individual target locations. Using the centroid of the receptive fields as the fixation point, the fixation has been shifted,

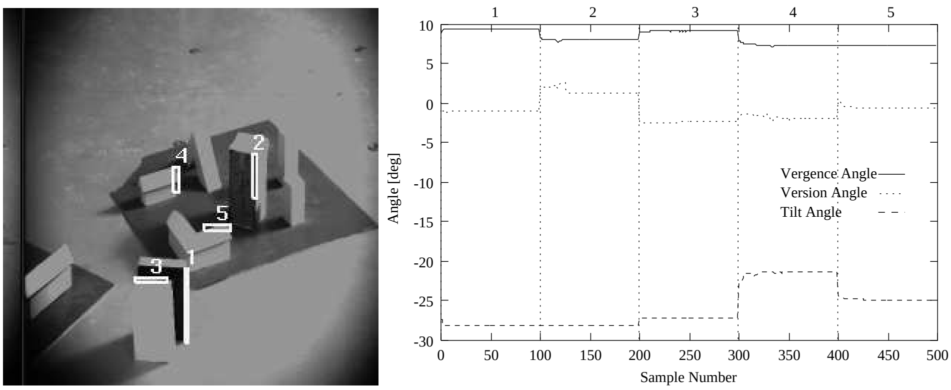

Fig. 7.: **Left:** The receptive fields (white) on top of the input image. The numbers denote the rank in which they were selected, with number 1 as the most salient region. **Right:** shows the Vergence, Version and Tilt angles during fixation shift. Shifts occur at sample 0,100,200,300 and 400.

resulting in vergence-version-tilt angle changes as shown to the right in figure 7.

The control is performed by disabling all image processes, while changing fixation. These processes are engaged again when the fixation change has been completed. This corresponds to that sampling is disabled during attentional shifts, thus the figure does not display any samples for the shift itself, only the fine tuning is shown. The attentional shifts are initiated every 100 iterations, thus the visual processes have 100 iterations to fine tune the fixated position. As it may be seen from the plot, the tuning only lasts about 10-15 samples. The combined vergence and version control fixates all of the targets to be within a 2-4 pixel distance from the image center, which is satisfactory. The tilt control on the other hand shows some severe problems, particularly when shifting attention from location 1 to 2 (around sample 100), where a significant tilt change is required. The main reason is related to the high amount of backlash in the tilt motor (about 1°), which is particularly pronounced in the region the head is working in during this experiment. When the head tilts slightly back and forwards from the initial position, the entire rig tilts about a degree, which is a problem the vision system does not compensate for. Though for the other attentional shifts, the effect is less pronounced, as seen when shifting from 4 to 5 at sample 400.

5 Summary

A general low level system design has been proposed for control of camera heads in a fixation based scheme using equal eye dominance control. The proposed system contains elements for fixation, tracking and attention selection/shifting, which is accomplished by separate control of the Cyclopean angles "version", "vergence", and "tilt". An integration strategy for fusing multiple cues for control of the vergence angle has also been proposed.

The designed system architecture has been evaluated through an implementation and with subsequent experiments, using the Aalborg head. This head has 4 rotoric degrees of freedom (2 vergence motors, plus common pan and tilt motors), and 6 optical degrees of freedom (focus distance, aperture, and focal length, in either lens).

The main results obtained are:

- The fusion of two range cues allows for full 3D tracking even when no disparity match is obtained as accommodation ranging is monocular. While this also could be facilitated within leading eye tracking systems, the separation of range and position in the equal eye dominance architecture makes the fusion process inherently simple.
- Experiments involving attention selection from a wide variety of scenes have been performed, and the system has been shown capable of fixating both static and dynamic objects as both data and model driven responses.

– Based on the experience obtained with the designed and implemented system it is asserted that the proposed architecture is an efficient way of organizing the low level vision system for a fixation based system.

References

1. Yiannis Aloimonos, Isaac Weiss, and Amit Bandyopadhyay. Active vision. In *Proceedings of the First International Conference on Computer Vision*, pages 35–54, London, UK, June 1987. IEEE, IEEE Press.
2. Erik Paul Krotkov. *Active Computer Vision by Cooperative Focus and Stereo.* Springer Verlag, New York, 1989.
3. Dana H. Ballard and Christopher M. Brown. Principles of animate vision. *Computer Vision Graphics and Image Processing*, 56(1):1–21, July 1992.
4. Dana H. Ballard. Animate vision. *Artificial Intelligence*, 48:57–86, 1991.
5. C. M. Brown, D. H. Ballard, T. Becker, R. Gans, N. Martin, T. J. Olson, R. Potter, R. Rimey, D. Tilley, and S. Whitehead. The Rochester Robot. Technical Report 257, University of Rochester, Computer Science Dept., 1988.
6. David J. Coombs and Christopher M. Brown. Cooperative gaze holding in binocular vision. *IEEE Control Systems*, pages 24–33, June 1991.
7. H. I. Christensen, K. W. Bowyer, and H. Bunke, editors. *Active Robot Vision.* Series on Machine Perception and Artificial Intelligence. World Scientific, 1993.
8. A. L. Yarbus. *Eye Movements and Vision.* Plenum Press, 1967.
9. D. A. Robinson. The oculomotor control system: A review. *Proceedings of the IEEE*, 56:1032–1049, 1968.
10. James L. Crowley and Henrik I. Christensen, editors. *Vision As Process.* ESPRIT Basic Research Series. Spring Verlag, 1995.
11. D. A. Robinson. Why visuomotor systems don't like negative feedback and how they avoid it. In Michael A. Arbib and Allen R. Hanson, editors, *Vision, Brain, and Cooperative Computation*, chapter 1, pages 89–107. MIT Press, Cambridge, MA, USA, second edition, 1988.
12. R.H.S Carpenter. *Movements of the Eyes.* Pion, London, 2nd. edition, 1988.
13. James J. Clark and Nicola J. Ferrier. Modal control of an attentive vision system. In *Proceedings of the Second International Conference on Computer Vision*, pages 514–523, Tampa, Florida, December 1988. IEEE, IEEE Press.
14. Kourosh Pahlavan, Thomas Uhlin, and Jan-Olof Eklundh. Dynamic fixation. In *Proceedings of the Fourth International Conference on Computer Vision*, pages 412–419, Berlin, Germany, 1993. IEEE, IEEE Press.
15. Claus Siggaard Andersen and Henrik Iskov Christensen. Using multiple cues for controlling an agile camera head. In *IAPR Workshop on Visual Behaviors*, pages 97–101, Seattle, Washington, USA, June 1994. IAPR, IEEE Press.
16. Claus Siggaard Andersen. *A Framework for Control of a Camera Head.* PhD thesis, Laboratory of Image Analysis, Aalborg University, Denmark, January 1996.
17. Claus Siggaard Andersen, Jan Juul Sørensen, and Henrik Iskov Christensen. An analysis of three depth recovery techniques. In P. Johansen and S. Olsen, editors, *Proceedings of the Seventh Scandinavian Conference on Image Analysis*, pages 66–77, Aalborg, Denmark, August 1991. IAPR, Pattern Recognition Society of Denmark.
18. Jun Ma and Søren I. Olsen. Depth from zooming. Technical Report 88/14, Dept. of Computer Science, University of Copenhagen, Copenhagen, Denmark, 1988.
19. Sean M. Culhane and John K. Tsotsos. An attentional prototype for early vision. In Giulio Sandini, editor, *Proceedings of the Second European Conference on Computer Vision*, volume 1, pages 551–560, Santa Margherita Ligure. Italy, May 1992. Springer Verlag.

This article was processed using the TEX macro package with SIRS96 style