

An Architecture for a VLSI Sensory-Motor System for Obstacle Avoidance

David Claveau and Chunyan Wang
Electrical and Computer Engineering
Concordia University
Montreal, Canada

Corresponding Author:

Chunyan Wang
chunyan@ece.concordia.ca
(514) 848-3120 (voice)
(514) 848-2802 (fax)

Electrical and Computer Engineering
1455 de Maisonneuve, West, S-H-961
Montreal, Quebec H3G 1M8
CANADA

Abstract

This paper presents a signal processing architecture for a sensory-motor system based on the smart sensor paradigm. The architecture is designed for an obstacle avoidance task by a mobile robot in an unstructured environment. Drawing inspiration from the field of behavior-based robotics, the development of the architecture is guided by an emphasis on the requirements of an obstacle avoidance behavior for a mobile robot. The architecture is simple enough for a smart sensor, but incorporates features which enable it to deal with realistic, unstructured environments. It differs from existing systems by using a special foveation scheme to facilitate the detection of real-world objects. From this, a motor control signal is produced by using a biologically inspired technique of aligning sensory and motor maps. The effectiveness of the architecture is explored through computer simulation, including an obstacle avoidance simulation in a 3D virtual environment.

Keywords: obstacle avoidance, autonomous robots, smart sensors

Autonomous robots which are meant to operate in remote or hazardous places will often need to perform visually guided behaviors in unstructured environments. Obstacle avoidance is an example of such a behavior. Traditional approaches to endowing robots with visually guided behaviors make use of conventional vision systems based on the combination of a camera and a digital processor. They place emphasis on constructing detailed models of the environment, from which appropriate actions are inferred. The processing is complex and computationally intensive and often leads to systems which are slow to react. The algorithms are also difficult to integrate into a monolithic system. Some of the traditional techniques to construct models, like stereo vision [1] and optical flow [2], also place strong constraints on the environment, such as the need for adequate texture or detailed features which the algorithms require to determine the depth of objects in the scene.

In contrast to traditional systems using general-purpose computers or DSPs, smart sensors integrate optical sensing and signal processing at the pixel level. They have been shown to be successful in applications such as visual tracking [3][4]. For an obstacle avoidance task for autonomous robots, a VLSI system should be able to detect obstacles which have irregular shapes and complex textures in order to deal with unstructured environments. It should also be easy to interface to actuators without the need for extra microcontrollers to decipher its output. To achieve such a VLSI sensory-motor system, new architectures and strategies are needed.

Two common approaches to achieving sensory-motor functionality are to use correlation-based motion detectors [5][6] and gradient-based position trackers [7][8][9]. An example of the latter is a complete VLSI sensory-motor system reported by Maris and Mahowald [7]. Its architecture consists of a contrast sensitive retina followed by a winner-take-all operation to detect the location of the pixels with highest contrast. The system exemplifies the efficiency of pixel-level processing but is limited to detecting clear and consistent edges in a scene rather than complex, natural objects. Also, the chip area required for the operations is quite large.

Another example of a smart sensor that supports visually guided behavior was reported by Etienne-Cummings [6]. It makes use of a two level foveation scheme which allows both tracking and acquisition to be implemented on the same focal-plane. Similar to the first example, pixel-level edge detection is used, limiting the system to track or avoid only simple objects with well defined contours.

The sensory-motor architecture proposed in this paper uses a special foveation scheme to facilitate the detection of real-world objects for an efficient obstacle avoidance behavior in an unstructured environment. The basic concept and the structure of this architecture, as well as the simulation results will be presented in the following sections.

2.1. Obstacle avoidance behavior

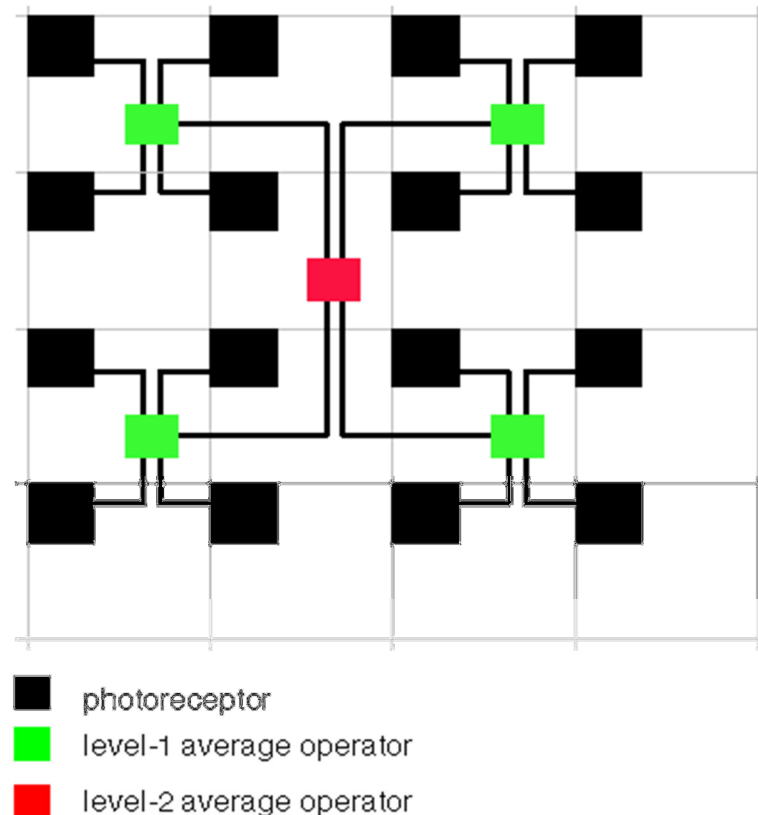
Inspired by the behavior-based approach to robotic control [10][11][12], the development of the architecture begins with a consideration of the required behavior. For an obstacle avoidance behavior, a robot should steer its body away from any stimulus which might signal that a significant object is in its pathway. As it moves, it should be orienting itself towards an open part of the pathway in front of it. The basic requirements are:

- 1) Real-world obstacles must be detected and located relative to the robot's body,
- 2) The robot should steer to the left or right around the obstacle in a reactive manner.

If, as the first requirement suggests, only relative locations are needed, then actual positions in a coordinate frame need not be computed. The first requirement also suggests that obstacles do not need to be recognized or even resolved with much precision; they only need to be detected. In addition, resolving power does not even have to be the same for all parts of the scene. More distant parts of the scene do not need high resolution because in the distance only the detection of large objects is needed to make steering decisions. Space-variant or foveated sensing schemes have been shown to have these qualities, in both biological [13] and engineered systems [14]. Also, foveation has been shown to be useful in creating a frame of reference around a point in a scene, the point of fixation [15]. This 'fixation frame' can facilitate the localization of objects relative to the point of fixation. Thus, the first architectural decision for the sensory-motor system in this paper is that a foveated sensing scheme should be used. The particular scheme proposed in this paper is described below in section 2.2.

The second requirement of the obstacle avoidance behavior suggests that for a typical wheeled robot, only a left/right orientation decision is needed to avoid an obstacle. This greatly simplifies the task of converting the visual signal to a motor signal. The fact that only a left/right decision needs to be taken can be exploited by aligning the map of the visual field (as represented on the sensor), with the center of the robot's body. If more obstacles are on the left of the visual field, then they are to the left of the robot's body and the robot should turn to the right. This is a simple alignment of a sensory map and a motor map, and is the second organizing principle for the architecture. An enhanced version of this alignment is described below in section 2.3.

There have been several sensor architectures reported in the literature which make use of foveation schemes [14][16]. All of these schemes vary the sizes and the placement of individual receptors in the array, e.g. log polar arrangements. However, for a VLSI smart sensor system, an architecture based on a homogeneous array of uniform pixels is easier to design and extend. The pixel can be designed like a standard cell with a constant pitch and can be arranged to abut with common power rails and interconnections. This would reduce the number of unique cells that must be designed and could make it easier to scale or modify the system. Thus, a foveation scheme based on a homogeneous array of uniform receptors is proposed. In this scheme, called post-receptor foveation [17], the variation in resolution is achieved by spatially averaging the signals from the photoreceptors. The basic structure is shown in Fig. 1. Each photoreceptor produces a signal, a current or a voltage, which is proportional to the intensity of the light which projects onto its area from the scene. The average of the signals from a 'field' of four adjacent photoreceptors is called a level-1 average. Following this, the average of four adjacent level-1 averages is called a level-2 average and so on, in a hierarchical manner. In this way, each level defines a set of fields which cover the array of photoreceptors.



A larger field contains or overlaps the smaller fields which contribute to its average value. The hierarchical simplicity of the scheme resembles the pyramid data structures often used in image processing [18].

In the post-receptor foveation scheme, the available resolutions are determined by the interconnection structure of the array. The resolution that can be obtained by accessing the level-1 fields is four times the resolution that can be obtained by accessing the level-2 fields. A fovea can be defined over an area where the smallest fields (eg. level-1) are used for processing. In this way, the fovea can be placed anywhere on the array and can assume different shapes. This is similar to the situation in some birds that are known to have an elongated fovea (to align with the horizon), or even multiple foveas as in the case of the owl [19]. These particular configurations of the fovea help to extract important information from the environment. Using a homogeneous array as a substrate, a diagram of a configuration that may be advantageous for obstacle avoidance with a small wheeled robot is shown in Fig. 2 (not showing the average operators and interconnections). The fovea is at the bottom of a 32x16 array. This can be used to create a frame of reference centered at the foot of the robot. Small obstacles which are close will be imaged with finer resolution than distant objects which will only appear if they are large and prominent.

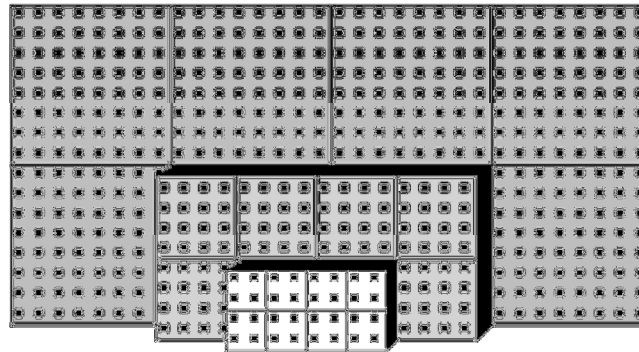


Fig. 2. Field sizes for a post-receptor foveation of a homogeneous array designed to facilitate obstacle avoidance for a small robot.

The post-receptor foveation offers some potential advantages when it is applied in realistic, unstructured environments. The hierarchical averaging produces a type of low-pass filtering of the image which is increasingly coarse away from the fovea. This can make complex objects like bushes or rocks appear more like a block of uniform intensity, facilitating their detection based on intensity differences. Also, emphasizing parts of the visual

distributed evenly across the array.

2.3. Sensory-motor map alignment

Avoiding obstacles is something that humans and animals do effortlessly. Regions of the brain involved in sensory- motor transformations are thought to use overlapping and aligned ‘maps’ of the visual field and of the body’s motor control system [20]. This strategy is the inspiration for a very simple scheme applicable to typical wheeled robots. An example configuration for an autonomous robot with a VLSI sensory-motor system is shown in Fig. 3. The sensory- motor system is mounted on the front of the robot with the center of the sensor aligned with the midline of the body. Making use of the strategy in [21], the sensor is angled toward the pathway so that the point of fixation is immediately in front of the robot where the path is assumed to be ‘open’. If all potential obstacles are on the ground plane (or near to it), then the more distant they are from the robot, the higher up they will appear on the image plane. Obstacles can then be located relative to the open path in front of the robot by using this measure of depth along with a determination of whether they lie more to the left or right of the point of fixation. Since the center of the sensor is aligned with the midline of the robot’s body, the point of fixation and the visual field are also aligned with the body. This means that if an obstacle lies more on the left side of the sensor, then the robot should steer to the right. By using spatial derivatives and other visual cues to determine where the open pathway is limited by obstacles, a continuous signal from each side of the sensor can be produced, the magnitude of which indicates the degree to which the side is obstructed. Thus, a left/right control signal can be obtained through this simplified transformation. Fig. 4 shows an example 16x8 sensor array configured in this manner. The fovea is at the bottom of the array in order to establish a fixation frame centered on the open pathway immediately in front of the robot. This allows the rest of the height of the sensor to be used to map scene depth. The extent of the open pathway on each side is expressed as a continuous left and right control signal, $c_L(t)$ and $c_R(t)$.




Fig. 3. Top view of a simple autonomous robot with a VLSI sensory-motor system.

The simple association of the left and right halves of a sensor with the left and right halves of the robot's body and the visual field can be further enhanced. Central regions of the sensor can be associated with the 'nose' of the robot and peripheral regions can be associated with the sides. This refinement can be used to add a weighting scheme to the obstacle detection cue such that objects which are more immediate and closer to the nose will elicit a stronger signal for the robot to turn away from them. Central, diagonal and side (or peripheral) regions may be distinguished to refine the response of the system. A spatial derivative across the smaller fields can have a larger weight, or gain factor, than for the next level field because it would indicate a more immediate obstacle. In the next section, a description of a complete architecture shows more clearly how these weights are used.

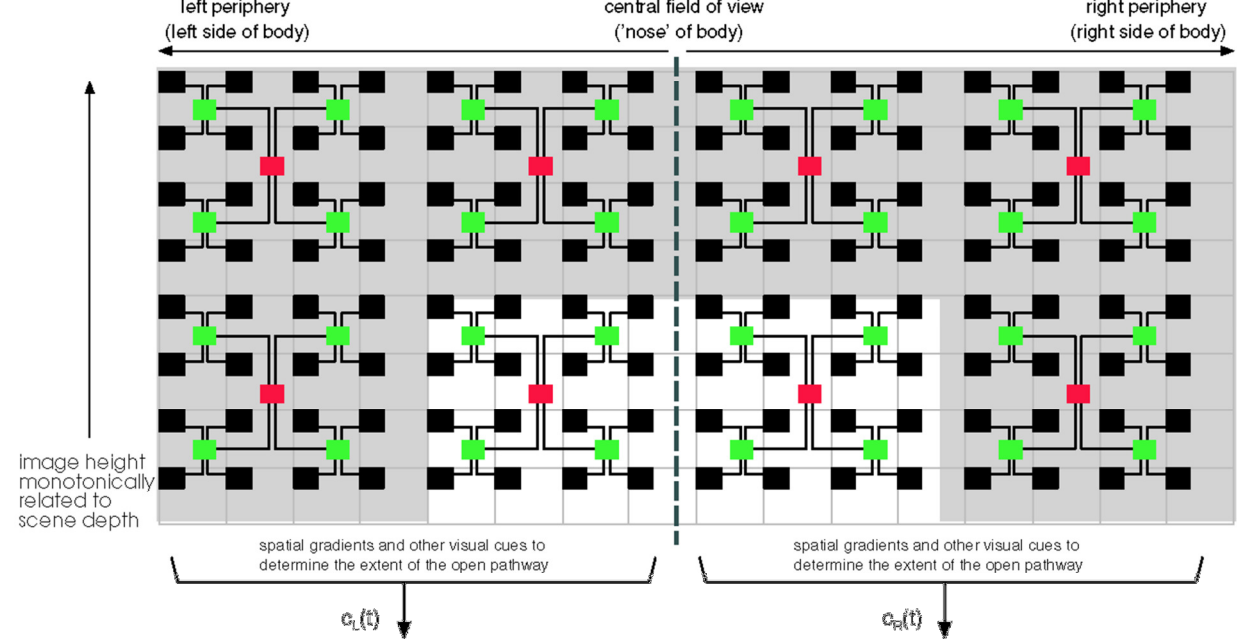


Fig. 4. An example 16x8 homogeneous sensor array divided into left and right halves. The fovea, shown above as the unshaded region, is aligned with the midline of the body to enable a left/right steering decision.

2.4. Architecture of the sensory-motor system

The post-receptor foveation scheme described in §2.2 and the enhanced sensory-motor map alignment scheme described in §2.3 can be applied together to develop a sensory-motor system for obstacle avoidance. The architecture is shown in Fig. 5. For clarity, only the right side of the system is shown. Comparators are used to find the spatial derivatives for the three regions: center, diagonal and side. The center derivative can signal an object in the robot's face, while diagonal and side derivatives are progressively less frontal. The derivatives should be weighted differently to reflect this. For example, the second level weight factor for the center is shown as K_{2C} and for the diagonal and side as K_{2D} and K_{2S} . In most cases they should follow the general relationship: $K_{2C} > K_{2D} > K_{2S}$. These factors can be related to the shape of the robot's body (narrow/wide) in order to determine more precise ratios. Derivatives for lower levels will signify more immediate obstacles so that if K_{iX} is used to refer to all gain factors at level i , then it should be that $K_{iX} > K_{(i+1)X}$, for $i \geq 1$. The angle at which the sensor is mounted on the robot (the vertical gaze angle) can be used to determine more precise inter-level ratios for these constants. Finally, by summing the scaled derivatives a continuous signal is produced to indicate the extent of the free pathway. The difference between left and right signals can then be used to steer the robot.

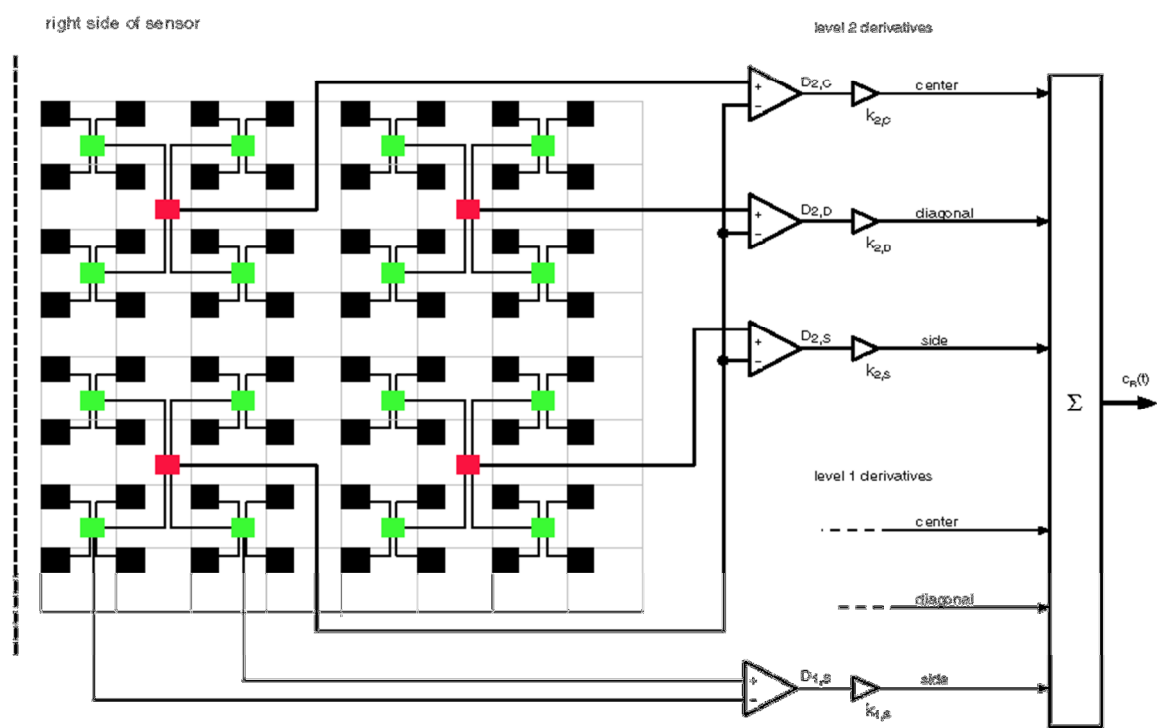



Fig. 5. The right side of a symmetric architecture. Hierarchical pixel interconnections are used to compute averages, spatial differences, weights, and a summation of all levels. Only one of the three level-1 comparators is shown. Total array size would be 16x8 with two levels of foveation ($L = 2$).

3. Simulation

The architecture was simulated to evaluate its effectiveness in realistic environments. Two simulations were carried out. The first used a still image and the second used a dynamic virtual environment. Both were based on the mathematical model described in the following sub-section.

3.1. Mathematical model

Guided by the post-receptor foveation scheme of §2.2, the mathematical model used to describe the architecture is very simple. Fig. 6 shows the right half of a homogeneous array of photoreceptors divided into fields according to the post-receptor foveation scheme. The fields are organized into levels where each level consists of four fields. Each field belongs to one of the regions: center, diagonal, side, and 'path'. The value of each field is the average of the values of its constituent pixels.



- f_{2c} : signal from level- n center field
- f_{2d} : signal from level- n diagonal field
- f_{2s} : signal from level- n side field
- f_{1p} : signal from level- n path field

Fig. 6. Right half of the array showing the notation used to refer to the four fields at each level where w_2 and h_2 are the width and height in pixels of the fields at level 2.

Using the notation shown in the figure, this can be expressed as the summation of the intensity values for all of the field's pixels, divided by the number of pixels in the field. If $I(x,y,t)$ denotes the intensity value of the pixel at position x,y and at time t , then the expressions for the center, diagonal, and side fields at the n th level are:

$$f_{nc}(t) = \frac{1}{w_n \cdot h_n} \cdot \sum_{x=1}^{w_n} \left[\sum_{y=h_n}^{2h_n} I(x, y, t) \right], \quad n \geq 1 \quad (1)$$

$$f_{nd}(t) = \frac{1}{w_n \cdot h_n} \cdot \sum_{x=w_n}^{2w_n} \left[\sum_{y=h_n}^{2h_n} I(x, y, t) \right], \quad n \geq 1 \quad (2)$$

$$f_{ns}(t) = \frac{1}{w_n \cdot h_n} \cdot \sum_{x=w_n}^{2w_n} \left[\sum_{y=1}^{h_n} I(x, y, t) \right], \quad n \geq 1 \quad (3)$$

where w_n and h_n are the width and height in pixels of the fields at level n . For the 'path' field, at level-1 the

$$J_{1p}(t) = \frac{1}{w_1 \cdot h_1} \cdot \sum_{x=1} \left[\sum_{y=1} I(x, y, t) \right], \quad (4)$$

and for subsequent levels, the values for the four fields from the previous level can be used to simplify the computation. Using these expressions, the right control signal, $c_R(t)$, would be:

$$c_R(t) = \sum_{n=1}^L K_{nc} \cdot |f_{nc}(t) - f_{np}(t)| + \sum_{n=1}^L K_{nd} \cdot |f_{nd}(t) - f_{np}(t)| + \sum_{n=1}^L K_{ns} \cdot |f_{ns}(t) - f_{np}(t)| \quad (5)$$

where L is the highest level.

For the purposes of a complete obstacle avoidance simulation, the control signals were used to determine the change in a robot's orientation and forward motion. The change in orientation angle, denoted $\Delta a(t)$, is proportional to the difference between the left and right control signals, with a scaling factor K_a , to convert to an angular measure like radians:


$$\Delta a(t) = K_a \cdot [c_L(t) - c_R(t)] \quad (\text{radians}) \quad (6)$$

The forward motion of the robot in the dynamic simulation was fixed at 30mm/sec. as in [12].


3.2. Simulation with a still image

First, a simulation was performed to determine if it was possible to derive reasonable control signals, $c_R(t)$ and $c_L(t)$, from a still image of a real scene. The test image was that of a black cube on a textured surface as shown in Fig. 7 (a). The image was subsampled down to 16x8 to match the architecture in Fig. 5, and to show that the architecture can be effective at such a low resolution. The post-receptor foveation algorithm was applied to the image. The image after the foveation step is shown in Fig. 7 (b). For each level of foveation, the differences between the path field and the surrounding center, diagonal and side fields are determined, as in equation (5), but with unity


unambiguously. The results show that a clear decision to turn to the right can be derived from the gradient map. They also show that the patterns in the scene were not a distraction to forming a meaningful gradient map. Nor were they necessary for the detection of the obstacle.



(a) Scene with simple obstacle on textured surface



(b) Foveation applied to the above image



(c) Gradient map for the foveated image above

Fig. 7. Test image and simulation results indicating an obstacle on the left.

For the second simulation, a virtual environment was created to simulate an unstructured environment in order to test the obstacle avoidance performance of the architecture in a more realistic way. The program was developed using the OpenGL graphics library. Its user interface is shown in Fig. 8. The view of the virtual environment was designed to represent how it would appear from the view of a sensory-motor system mounted on the front of a robot, as shown in the top left panel of the interface. The environment was made to look realistic by applying textures to the floor and other objects. The top right panel shows a top view of the room in which each dark shape is one of the obstacles shown in the perspective view. The bottom left panel shows the foveated view. Two types of obstacles were used. One was a plant-like structure designed to emulate obstacles which are difficult to deal with using traditional techniques. For example, infrared sensors would have difficulty detecting such an obstacle because of the complex reflectance patterns. Also, classical edge detection may not be efficient because of the shape of the obstacle and the pattern on the floor. It would lead to a large amount of data which would have to be sorted through.




Fig. 8. User interface for the virtual environment simulator. The top left panel shows the view of the room from the perspective of the robot. The bottom left panel shows the foveated image. The top right panel shows a top view of the room. The dark shapes are the plant-like obstacles in the room. A control panel is at bottom right.

(a). It was designed to represent the smallest obstacle that must be avoided. In the scale of the simulator the hole measured 10cm in diameter and the plant-like obstacles were approximately 50-100 cm in diameter and 50-100 cm in height. Assumptions implicit in the simulation include the absence of shadows, and the use of a flat surface such as an office floor.

A simulation can be run in such a way that the view of the robot is updated frame by frame. Fig. 9 shows a sequence of frames produced by the simulator. Two obstacles, a hole that is very close and a plant which is further away and to the left, are used to show the obstacle avoidance behavior. The current view from the robot, illustrated in the top half of each image, shows how the robot should first move to the left to avoid the more immediate 'hole' obstacle. The foveated view shows that in the absence of a weighting scheme, a correct steering decision would be in doubt. With weights of:

$$\begin{array}{lll} K_{1C} = 4.0 & K_{2C} = 2.5 & K_{3C} = 1.0 \\ K_{1D} = 3.5 & K_{2D} = 2.0 & K_{3D} = 0.5 \\ K_{1S} = 3.0 & K_{2S} = 1.5 & K_{3S} = 0.5 \end{array}$$

(where the numeric subscript indicates the level and the 'C' is for center, 'D' for diagonal, 'S' for side) the closer object automatically has a greater effect and the simulated robot moves to the left as seen in the next two frames in (b) and (c).




Fig. 9. A sequence of frames produced by the simulator as two obstacles are encountered. The foveated image is below each frame. (a) Frame 1: two obstacles in view, (b) Frame 2: robot correctly steers to left of closer obstacle, (c) Frame 3: path is clear of obstacles.

A set of four simulation runs are shown in Fig. 10. The results for each simulation run were taken from the top-view of the room which showed the path of the robot as it moved around the obstacles. The first three runs, (a)-(c), were performed with a floor made to resemble a typical hardwood pattern. The fourth run, shown in (d), was performed on a floor of uniform color without texture. The black line represents the center of the robot's path and the grey portion represents the width of the robot's body, which was set to 50cm. In each case the robot avoided the obstacles. The results from the simulation runs demonstrated that the architecture leads to an effective obstacle avoidance behavior of the robot in an unstructured environment. They showed that it was effective with or without textures in the scene, which is an advantage over the approaches based on stereo or optical flow that depend on the presence of details. They also showed that the post-receptor foveation and the weighting scheme make it possible to appropriately avoid the conflict of controls in the case of multiple obstacles by prioritizing them based on their place in the visual field and their position with respect to the body of the robot.




Fig. 10. Simulation results: (a)-(c) show the results for a textured floor, and (d) shows the results for an untextured floor.

4. Conclusion

An architecture for a VLSI sensory-motor system has been proposed. Drawing on inspiration from biology and behavior-based robotics, it leads to an efficient obstacle avoidance behavior for a robot in an unstructured environment. The architecture includes a post-receptor foveation scheme which permits the efficient localization of obstacles through spatial derivatives. The foveation structure is hierarchical and is based on a homogeneous sensor array, facilitating its design and increasing its flexibility. It simplifies obstacle detection by producing a type of low-pass filtering specifically for obstacle avoidance, which makes complex objects (like bushes or rocks) appear like a block of uniform intensity. In addition, a weighting scheme is used to prioritize obstacles based on their place in the visual field. A left/right steering signal is produced by correlating a map of the visual field with a map which

simulations with a virtual environment. From an implementation point of view, the architecture includes only computationally simple operations which generally have simple implementations with a small number of transistors. These factors could lead to implementations with pixels that devote most of their area to the photoreceptor rather than to signal processing circuitry. This can improve the overall quality of the image signal. Future work will involve an ASIC implementation of the architecture. It also may involve exploring the inherent flexibility of the architecture with regard to the size and placement of the fovea and the values of the weights for different applications.

References

- [1] D. Murray, C. Jennings, "Stereo vision based mapping and navigation for mobile robots," presented at IEEE International Conference on Robotics and Automation, Albuquerque, NM, April 20-25, 1997.
- [2] D. Coombs, M. Herman, T. Hong, and M. Nashman, "Real-time obstacle avoidance using central flow divergence and peripheral flow," presented at the Fifth International Conference on Computer Vision, Cambridge, Massachusetts, June, 1995.
- [3] V. Brajovic and T. Kanade, "A VLSI sorting image sensor: global massively parallel intensity-to-time processing for low-latency, adaptive vision," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 1, pp. 67-75, February 1999.
- [4] T.G. Morris, T.K. Horiuchi, S.P. DeWeerth, "Object-based selection within an analog VLSI visual attention system," *IEEE Transactions on Circuits and Systems, Part II*, 45(12):1564--1572, 1998.
- [5] R.R. Harrison and C. Koch, "A neuromorphic visual motion sensor for real-world robots," presented at the Workshop on Defining the Future of Biomorphic Robotics, IROS 1998, Victoria, B.C., Canada.
- [6] R. Etienne-Cummings, J. Van der Spiegel, P. Mueller and M.Z. Zhang "A foveated silicon retina for two-dimensional tracking", *IEEE Trans. Circuits and Systems II*, Vol. 47, pp. 504-517, June 2000.
- [7] M. Maris and M. Mahowald, "Neuromorphic sensory-motor mobile robot controller with attention mechanism", *Neuromorphic Systems: Engineering Silicon from Neurobiology* (L. S. Smith & A. Hamilton, eds.), World Scientific, 1998.
- [8] G. Indiveri, "Neuromorphic analog VLSI sensor for visual tracking: circuits and application examples", *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, Vol. 46, Num. 11, pp. 1337--1347, November 1999.
- [9] T. Horiuchi and C. Koch, "Analog VLSI-based modeling of the primate oculomotor system," *Neural Computation*, 11(1), pp. 243-265, 1999.
- [10] R.A. Brooks, "New Approaches to Robotics," *Science*, 253, pp. 1227-1232, 1991.
- [11] R. C. Arkin, *Behavior-Based Robotics*. MIT Press, Cambridge, MA, 1998.
- [12] A. Duchon and W.H. Warren, "Robot navigation from a gibsonian viewpoint", in *Proc. IEEE International Conference on Systems, Man and Cybernetics*, Los Alamitos, CA, 1994, pp 2272-2277.
- [13] E.L. Schwartz, D.N. Greve, G. Bonmassar : "Space-variant active vision: definition, overview and examples", *Neural Networks*, 8, 1297-1308 (1995).

- [15] D. Ballard, "Animate vision", *Artificial Intelligence*, Vol 48, No. 1, pp. 1-27, February 1991.
- [16] R. Wodnicki, G. Roberts, and M. Levine, "A foveated image sensor in standard CMOS technology," in *Proc. Custom Intergrated Circuits Conf.*, 1995, pp. 357-360.
- [17] D. Claveau and C. Wang, "An architecture for a VLSI sensory-motor system for autonomous robots", IEEE/RSJ International Conference on Intelligent Robots and Systems, Switzerland, 2002.
- [18] S. Tanimoto and T. Pavlidis, "A hierarchical data structure for picture processing," *Computer Graphics and Image Processing*, 4: 104-119, 1975.
- [19] Butler, A.B. and Hodos, W., *Comparative Vertebrate Neuroanatomy: Evolution and Adaptation*. New York: Wiley-Liss, 1996.
- [20] Stein, B.E. and Meredith, M.A. *The Merging of the Senses*, Cambridge, MA: The MIT Press, 1993.
- [21] I. Horswill, Specialization of Perceptual Processes, Ph. D. Thesis, Massachusetts Institute of Technology, May 1993.