Active Sensing at a Microscopic Scale

James J. Clark
Robert P. Hewes

Division of Applied Sciences
Harvard University
Cambridge, MA

Abstract

We examine the possibilities afforded by the application of active sensing principles at a microscopic scale.

Active point sensor systems consist of collections of simple sensing elements that have motor, computational, and communication capabilities. Depending on the laws of motion programmed into these elements, complex operations can be performed, as diverse as nonlinear diffusion, shape from shading, and deformable template matching.

Active point sensors systems represent an entirely new way of implementing sensing devices. In contrast to sensor arrays made using present techniques, they permit the integration of complex information processing on the same substrate as the sensing elements, and can be adaptive, altering their processing based on changes in the state of individual sensing elements. Furthermore their output is asynchronous and object oriented, reducing communication bandwidths and easing interfacing to other systems.

1 Principles of Active Sensing

Active sensing can be defined as *a process that uses controlled motions of an observer to aid in the performance of sensory information processing tasks.* Recent research into active vision [1] and ecological optics [14] has made it clear that the ability to control the motion and associated processing of a visual sensor is extremely valuable in solving visual information processing problems. For other sensory modalities such as touch and smell, the connection between motion and sensing has been known for a much longer time. Effective tactile sensing absolutely requires that a tactile sensor be moved in a purposeful exploratory pattern. For example, determination of object texture by tactile means requires a movement of the tactile sensor across the object surface in a direction that maximizes the temporal response of the sensor (i.e. across the grain of the texture).

For descriptive purposes active sensing techniques can be categorized as employing one or more of the following basic principles: *Regularization of Sensory Processes; Exploration or Search for New Information; Adaption.* The first of these principles is that of the regularization of sensory processes. Many standard computational vision tasks are ill-posed in the sense that, given the input sensory data, there is either no unique solution to the sensory information task, or if there is a solution, the solution is unstable with respect to small perturbations in the input data. Aloimonos et al [1] show that some of these ill-posed vision tasks can be made well posed (and often linear as well) if extra information in the form of images taken at different times and from different viewpoints is available. Thus, by using active vision techniques one can regularize an ill-posed vision problem, where regularization refers to the mathematical process of making an ill-posed problem well-posed through the addition of a stabilizing process [19].

Perhaps the most widely used principle of active sensing is that of exploration, or search, in which sensor motions are performed in order to allow sensing of portions of the environment that were previously out of the range of the sensing elements. This exploration can be done in an unguided fashion, wherein the motions are arbitrary, or as guided motion wherein the sensor motions follow some strategy aimed at optimally extracting the desired information from the environment. The proposed active vision systems of Bajcsy [2], Burt [5] and Stansfield [23] are examples of exploratory systems that use navigation about the world in order to obtain information about the world. In Stansfield's application, for example, tactile exploration is used to obtain information about the shape of objects for object recognition. The purpose of the exploration in this case is to extract features to be used in the recognition process. Thus the exploration is guided by expectations as to where object features are located.

The third principle of active sensing that we consider important is that of (temporal) adaption. As a sensing system moves about the information it gains can be used to alter the form of the information processing that is done. This alteration is typically one of two types. In the first type, certain assumptions or constraints used in the sensory processing algorithms can be changed based on consistency measures. For example [8], the smoothness constraint used by some shape from shading algorithms cause a distortion or flattening of the measured shape of an object. As the imaging device moves about the space, and the shape from shading process is repeated, the perceived shape of the object will appear to deform and the object will appear non-rigid. The inconsistency of the shape over time can be used to adapt the smoothness constraint. The second type of adaption occurs when the goals of the information processing changes in response to the new sensor information obtained as the sensor moves. For example, an exploration phase may change into an object tracking phase, or a more complex object recognition process may be introduced once an "interesting" object is found by more coarse object detection processes.

2 Active Point Sensor Systems

In this paper we wish to introduce the idea of applying the above principles of active sensing to microscopic scale systems, with the intention of constructing a new class of sensing device. To date, active sensing principles have mainly been applied to robotic systems, whereby the motion of sensors such as video cameras and tactile sensors are carried out through the degrees of freedom of the robot. We want to extend this process to systems wherein collections of very simple, microscopic, point sensing elements having motion capability are used to implement a new class of macroscopic sensing devices, which replace the video cameras and tactile sensors currently used in robotic systems.

The basic building block of this new approach is the *active point sensor.* An active point sensor is an entity which consists of a single

This research was supported in part by the ARO Brown/Harvard/MIT Center for Intelligent Control Systems under grant DAA103-86-K-0171, by the Harvard/University of Maryland NSF Systems Research Center under grant number CDR-85-00108, and by DARPA grant N00039-88-C-0163.

sensing element, a motor system which allows it to move through space in some fashion, and a communication system which allows it to transfer information to and from other sensors or information processing units. It has simple computing capabilities which allow it to determine its motor behaviour based on its current state. The current state of the active point sensor depends on external inputs, the previous state, and the current measurement from the sensing element.

The task of an active point sensor is to take measurements with its sensor, move in some prescribed fashion based on the measured sense data, and conditionally, depending on its current state, transmit information to other units in the system. The laws which specify the dependence of the motion on the sense data and the active point sensor state can change or adapt as a function of the sense data and the system state. In general the overall behaviour of an active point sensor will be dictated by the goals of the controlling system of which the sensor forms a part.

The operation of active point sensor systems can be illustrated by examining an instance of such a system that occurs in nature. Current models of vertebrate immune system functioning (the clonal selection theory [4]) imply that the immune system contains an active sensing component. In this system antibodies act as sensors that can detect (by physical contact) foreign substances (antigens) in the body. When they detect the antigen specific to the antibody, they signal cells known as lymphocytes to create more antibodies of the same type. The motion of the sensing elements (antibodies) is random in this case (basically thermal diffusion) but the net effect is that the system accurately localizes the presence of antigens of a certain kind (sensory modality) by the presence of large numbers of antibodies to that antigen. The lymphocytes act as an adaptive mechanism to alter the dominant sensory modality of the immune system. Although the operation of the immune system may seem to be quite different than the operation of usual image processing systems, we will see below that we can devise image processing systems that are based on exactly the same principles as the immune system: mobile sensing elements (antibodies) capable of adaption of sensory modality (types of antigens detected) combined with communication with control units (lymphocytes).

3 Stochastic Processing By Active Point Sensors

As can be seen in the example of the immune system, the motion of active point sensors need not be purposeful in order for them to be useful information processing units. There is a wide range of sensory processing tasks for which random (thermal-like) motion of active point sensors will suffice. An example relevant to computer vision systems is that of nonlinear smoothing or edge enhancement via nonlinear diffusion. These image processing methods are currently of interest in computational vision research [21]. In the process of Perona and Malik [21] an image is input as an initial condition to a nonlinear diffusion equation which is then allowed to evolve. After a suitable length of time the solution to the diffusion equation is a new image that is a smoothed version of the inital image, except that, unlike standard linear smoothing algorithms, the edges, or regions of rapid change in image intensity, are not smoothed away. The equation used in the process of Perona and Malik [21] is:

$$I_t = c(x,y,t)\Delta I + \nabla c \cdot \nabla I \quad (1)$$

where I is the (smoothed) image, and $c(x,y,t)$ is a time and space dependent conductivity factor. If we let c be a function of the image as follows:

$$c(x,y,t) = \frac{C}{1+K|\nabla I|^2} \quad (2)$$

The conductivity is seen to be low when the image gradient is high. This prevents the image from being excessively smoothed near edges. Such a system can be implemented with a set of active point sensors moving in random thermal motion if we allow the freedom of specifying the "effective mass" or "mobility" of the point sensors to be a function of the sensed image gradient. The diffusion "constant" of the sensor's thermal motion can be identified with Perona and Malik's $c(x,y,t)$ function. Thus if we can adjust the sensor's mobility according to the equation for $c(x,y,t)$ given above we will be able to duplicate Perona and Malik's edge enhancement method with our active point sensors. Note that, in the active point sensor implementation, as time progresses the sensors will tend to cluster near edges, as these are regions of low diffusivity. Thus edges will be indicated by the location of high concentrations of active point sensing elements. The positions of the point sensors can then be communicated in some fashion to other processing elements.

The process of deriving the laws of motion for a general diffusive system as described above can be implemented using the methods of stochastic differential equations [11]. For example, suppose we wish to implement the following generalized diffusion equation:

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho f(x)) + \frac{1}{2}\nabla \cdot \nabla(\rho g(x)) \quad (3)$$

This equation can be seen to be a general form of a Fokker-Planck equation and hence we can find a stochastic system whose time evolution is a solution of it. In fact, if we have a system of point sensors whose motion is given by the following law:

$$dx = f(x)dt + g(x)dw \quad (4)$$

then the function ρ in the Fokker-Planck equation can be identified with the density of point sensors. The equation of motion is often referred to as an Ito or Langevin equation. The w in the Ito equation is a Wiener process and provides the stochastic behaviour. The first term in the stochastic differential equation [11] is a "drift" term, while the second is a "diffusion" term. By adjusting the form of the drift and diffusion functions the solutions to many different differential equations can be found by letting the point sensors move according to the Langevin equation and measuring the point sensor density ρ.

The function $g(x)$ acts as a diffusivity or conductivity. If this is non-uniform, the invariant measure (i.e. the steady state ρ), if it exists, will be non-uniform as well. If $g(x)$ is a function of sensory data, ρ will also. One can use this to peform certain sensory processing applications, such as edge enhancement. If $g(x)$ or $f(x)$ depends on $\rho(x,t)$, then the system will perform non-linear diffusion (such as the Perona-Malik operation), which can be used for more complex image processing than non-uniform diffusion. This approach requires that ρ be estimated, in order to determine the form of the motion laws. This estimate of ρ is termed the conditional density, as it is conditioned on the statistics of the path, $x(t)$.

The process of random thermal motion of active point sensors modulated by field measurements can be used to implement a wide range of useful image processing operations besides edge enhancement and image smoothing. Recent work by Yuille et al [26] has indicated that the task of obtaining the solution to the minimization of energy functionals can be converted into a task of solving an initial value partial differential equation. These PDE's are often similar in form to a nonlinear diffusion equation and can hence be solved by using the approach described above.

4 Non-Stochastic Active Point Sensor Systems

The motion of our active point sensors need not be random and, in general, will be purposeful, as their motion is intended to aid in the sensing process. For example, a tactile sensing element whose role is to find and describe ridges of objects would first look for a peak in pressure by moving along the pressure gradient direction and then, once it has found the peak, would move perpendicular to the pressure gradient in order to chart the extent of the object ridge. When it comes to the end of the ridge it should move away and look for another, possibly remembering where it was so it doesn't return.

In this section we describe a particular class of deterministic active point sensor systems; one that has surprising information processing capabilities.

There are classes of information processing problems which can be expressed in terms of solving first order nonlinear differential equations using the method of characteristics [13]. These equations can sometimes be put into the form of the Hamilton-Jacobi equation:

$$\frac{\partial J}{\partial t} + H\left(\frac{\partial J}{\partial x}, \frac{\partial J}{\partial y}, x, y, t\right) = 0 \tag{5}$$

where $H(p, q, x, y)$ is the Hamiltonian of the system. The equations of motion defined by the Hamiltonian are:

$$\frac{dx}{dt} = \frac{\partial H}{\partial p}, \frac{dy}{dt} = \frac{\partial H}{\partial q}, \frac{dp}{dt} = -\frac{\partial H}{\partial x}, \frac{dq}{dt} = -\frac{\partial H}{\partial y} \tag{6}$$

To solve a Hamiltonian based problem with an active point sensor system one would initialize the position of the sensors along some curve in the image on which $p(x, y)$ and $q(x, y)$ are known and let them move according to the laws of motion. At intervals they would transmit their estimates of $p(x, y)$ and $q(x, y)$ (updated using the differential equations for p and q given above). This approach is not restricted to Hamiltonian systems; many first order non-linear differential equations can be solved using the method of characteristics. Hamiltonian type systems are very common, however.

An example of the use of characteristic methods in image analysis occurs in shape-from-shading where, for a Lambertian surface (one whose reflectance function is given by $R(\hat{n}) = \hat{n} \cdot \vec{s}$ where $\hat{n}$ is the unit normal vector of the surface and $\vec{s}$ is the illumination vector) the object surface height function $u(x, y)$ is related to the image shading, $I(x, y)$ through the following Eikonal equation [3, 16]:

$$u_x^2 + u_y^2 = p^2(x, y) + q^2(x, y) = \left(\frac{1}{I^2(x, y)} - 1\right) = E^*(x, y) \tag{7}$$

This can be solved using the method of characteristics where we get the following equations of motion (that define the characteristics):

$$\frac{dx}{dt} = 2p(x, y), \frac{dy}{dt} = 2q(x, y) \tag{8}$$

and the update equation for the shape:

$$\frac{dp}{dt} = \frac{\partial E^*}{\partial x} = -\frac{\frac{dI(x,y)}{dx}}{I^3(x, y)}, \frac{dq}{dt} = \frac{\partial E^*}{\partial y} = -\frac{\frac{dI(x,y)}{dy}}{I^3(x, y)} \tag{9}$$

An active point sensor implementation of a shape from shading operation might function as follows: A set of active point sensors are initially distributed randomly about the image field. The sensors initially move to find a position where the p, q values are known. This is accomplished as follows. A given sensor assumes a value, perhaps randomly, for the pair (p, q). In general, this pair will not be the same as the actual (p, q) at the sensor's current location. As a first step to finding the point in the image where the actual (p, q) values are the same as the assumed ones, the sensor moves in such a way as to arrive at a point for which the constraint that $p^2 + q^2 = E^*(x, y)$ is satisfied. This could be done by a random walk, or with a spiral search pattern. Having reached such a point it will generally be the case that the actual (p, q) values are not equal to the assumed values. But we at least know that the E^* constraint is satisfied. The sensor, once it has reached this stage, alters it's motion law to become a characteristic follower, using the equations given above. It performs one step along the characteristic curve. If the assumed (p, q) values were equal to the actual (p, q) values at the starting point then the E^* constraint will still be satisfied after the step, otherwise it would not be satisfied. Thus we can check whether or not we have the correct (p, q) value. If it is incorrect we move along the iso-brightness contour (which will maintain the validity of the E^* constraint) a fixed distance and try another characteristic following step, and again check to see whether the E^* constraint holds. We continue this process until we get to a point where the E^* constraint holds after a charactistic following step. At this point we know that our assumed (p, q) pair is equal to the actual one. The sensor then continues following along the characteristic (which will be a function of the image) and periodically transmits it's current value of the p, q shape parameters.

5 Correlated Active Point Sensor Assemblies

Active point sensor assemblies are defined here as collections of active point sensors that are distinguished by close correlation of activity. The standard sorts of image sensor arrays, IR focal plane detector arrays, tactile sensor arrays and so forth can be thought of as examples of sensor assemblies. But these arrays use passive approaches to sensing and are hence limited in their capabilities, and must be paired with complex processing circuitry to obtain useful information from them. In addition these sensor arrays are inflexible in that they have a fixed spatial-temporal sampling pattern. Our vision of active point sensor assemblies is to develop sensor arrays which have none of these drawbacks. They are configurable, in that they can change sensory modalities and sampling topology, and they perform sensing in a purposeful and principled manner. These arrays produce only the information that is required for a given task and are silent otherwise. They move about in order to optimize the sensing process. The objective of active point sensor assemblies is to coordinate activity between active point sensors in order to extract spatially organized information which the active point sensors themselves cannot independently provide.

As an example of the use of active sensor assemblies to visual information processing, consider the "Snakes" paradigm of Kass, Witkin and Terzopoulos [17]. The basic "snake" model is a controlled continuity spline which is under the influence of forces derived from sense data (such as tactile or visual imagery). The elements that make up the snake move in a fashion that minimizes the following "energy":

$$E = \int_0^1 [E_{internal}(\vec{v}(s)) + E_{image}(\vec{v}(s))]ds \tag{10}$$

where the vector function $\vec{v}(s) = (x(s), y(s))$ is a parametrization along the arc length s of the snake. The internal energy $E_{internal}$ represents the correlation between the elements of the active sensor assembly (which are assumed to be located uniformly along the snake) and, in Kass et al's model, is given by:

$$E_{internal} = \alpha(s)\left|\frac{\partial \vec{v}(s)}{\partial s}\right|^2 + \beta(s)\left|\frac{\partial^2 \vec{v}(s)}{\partial s^2}\right|^2 \tag{11}$$

The first order term makes the snake act as though it was made of an elastic material, while the second order term makes the snake act like a thin metal plate. Adjusting the relative values of α and β allow one to tune the desired behaviour of the snake.

The image energy E_{image} represents the effect of the sense data on the shape of the snake. For example, if we wish for the snake to detect edges (by aligning with them) then we could use the following energy:

$$E_{image} = E_{edge} = -|\nabla I(x, y)|^2 \tag{12}$$

This will tend to make the snake move to regions of high image gradient value (i.e. edges). The internal energy term will act to keep the snake from becoming too irregular in shape, and essentially imposes smoothness on the snake's contour. Kass et al propose other image energy terms which allow snakes to be used for stereo and motion analysis.

Active point sensor assemblies can be made to implement snake based image processing algorithms. Local communication between neighboring elements in an active sensor assembly is maintained so that the internal energy can be computed. The elements in the sensor assembly move in directions that will jointly minimize the internal and image forces. The assembly can change shape in two ways: the component point sensors can move, or point sensors can be added or deleted to the assembly. The most straightforward approach is a Monte-Carlo-like method in which active point sensors are added and deleted to the assembly randomly (at a relatively low rate) and existing elements move thermally (randomly). If a given action (creation/deletion/thermal motion) results in a decrease in the snake's total energy then it is retained, otherwise the action is undone. In order to avoid local minima which could trap the snake, annealing must be performed using the artifice of randomly retaining a small proportion of the actions that increase the snakes total energy. The information regarding the shape of the snake could be transmitted in the form of a chain code [10] in which each element transmits its own part of the code, tagged in such a way that a receiver can sort out and reconstruct an number of such messages from differerent snakes.

The active point sensor assembly idea can also be used to implement deformable templates such as those described in [25] for face recognition. In the deformable template technique one tries to minimize an energy similar to that in the snake approach, but where the contributions to the internal energy function come from parametrized descriptions of the contours (e.g. parabolic sections). In an active sensor implementation of a deformable template, the active point sensors would move in ways which attempt to maintain the particular form of the template, while altering the parameters of the contour in ways which reduce the total internal energy of the contour.

6 Architecture of Integrated Active Sensor Systems

It is natural to consider how one might physically construct the active point sensor based systems described in the previous sections. Some ideas concerning important aspects of this implementation process are discussed below. We feel that the construction of micro-scale active sensing system provides a powerful alternative to the current approaches to the implementation of smart sensor chips (e.g. Mead and coworkers [15, 18, 20, 22] resistive grid based image sensing and processing chips).

A major aspect of the active sensor design process is determining how to get data out from the sensors and how to get commands regarding sensor behaviour into the sensors. In standard fixed, non adaptive, sensing arrays it is almost always possible to scan the array in an orderly fashion since the sensors do not change their position or function relative to each other. In active sensor systems, however, the component elements move about and do not maintain a fixed spatial relationship to each other. Furthermore the type of information produced by a given sensing element can be different from that of the others and may change over time as well. In light of this it seems evident that an asynchronous, object oriented method of transmitting information is indicated for active sensing systems. That is, information is output from the active sensor chip only when new information arrives, and the interpretation of the output data depends on the type of sensory object that is being output (e.g. visual edges, tactile texture, temperature gradients etc.).

A suitable communications method is one based on message passing techniques such as are used in computer networks (for example, EtherNet or the router on a Thinking Machines Corp. Connection Machine). In this approach, when a sensor decides that it has detected something some other unit should know about, it drops a message into a common transmission channel, tagged with information about the source and possibly the destination (or it might be broadcast to everyone by default). The output port of the sensing system then would be this common communications channel over which all sorts of different types of messages are being passed. Such a system is seen in the human body where biochemical messages (such as those used in the immune system) all flow through the blood stream, and are ignored by most systems and recognized only by those intended to receive the message. A typical computer vision sensor using this object oriented message passing technique would output a stream of data asynchronously, containing messages of various lengths and complexities; some would indicate the location of edge segments, some would describe the shape of snakes wrapped around regions of uniform texture, still others may signal the presence of objects moving upwards and to the left. A message passing based active sensor will clearly be more efficient than a set of fixed sensing systems that do the same job, as one does not have to add in extra circuitry to filter out irrelevant data (such as the velocity of points in the interior of a static object), since information is transmitted only when required.

In the basic form described earlier an active point sensor system would contain a number of active point sensors or active sensor assemblies. Each of these elements would contain a motor unit for the purpose of moving about in, and manipulating, its environment. In addition, the active sensor elements would contain sensing devices, limited computational capability, and communication circuitry. These systems can be constructed on a number of levels, from microscopic scale (such as the elements of the immune system, or Drexler's [9] molecular assemblers), to large scale (such as the attentive vision system for mobile robots that we have developed [6]). Each of these different scales will have their own particular technical challenges that must be faced in constructing such systems. The molecular scale systems of Drexler's [9] nanotechnology are beyond the current capabilities of current day technology. Work is proceeding in this area, however, and we should at least be thinking about how the principles embodied in our approach to active sensing can be put to use at molecular scales, and indeed, it is in this area that our work in micro-scale active sensing may eventually have its greatest impact. At the micron scale (i.e. at the scale of integrated solid state circuits) we have the technology for implementing sensing and computing elements (see, for example, [7, 20]). Motor technology at this scale, however, is in its infancy, and only a few primitive devices have been created [12, 24]. The outlook for the development of such devices is promising, however, and we should begin to design active sensing systems that take advantage of these microactuators. Until such microactuators are available we can *simulate* the motion of microscale sensors by implementing "virtual sensors" in a fixed array on an integrated circuit die. In the *virtual grid active sensor* chip "attributes" of a given active point sensor move through the chip via local transmission between neighboring virtual sensors. The actual sensors themselves do not move, just their "spirit" as it were. The attributes vector includes information such as current sensor modality, current velocity, and current sensor state (of its "program", if present).

7 Some Simulations

We have implemented a simulation of an active point sensor system on a SUN4 workstation, as well as on a MasPar 1024 node parallel computer. Figures 1 through 4 illustrate the action of an active point sensor nonuniform diffusion edge enhancement operation (i.e. the motion law of the sensors is given by equation 4 with $g(x) = 1/(1+K|\nabla I|^2)$). Figure 1 depicts the underlying image field I (what the sensors measure). Figure 2 shows the initial (uniform) distribution of the point sensors. The initial density of sensors is 0.2. Figure 3 shows the distribution of sensor elements after 50 time steps and figure 4 shows the distribution after 100 time steps. The dark pixels represent sensors for which the measured image gradient is greater than a certain threshold. The lighter pixels represent sensors for which the image gradient is less than the threshold. One could think of the sensors sending out a message (their position for example) only when their gradient value exceeds this threshold. In this way only relevant data, that of the position of edges, is transmitted. Note that, even after 100 time steps, only a portion of the image edges have been detected. This is due to the slowness of the thermal diffusion process. We can speed this process up by using the same sort of artifice as is employed by the human immune system. In this approach, whenever a sensor detects an edge (by exceeding the gradient threshold) it "creates" a new sensor nearby. The idea behind this is that edges will tend to be coherent in space, so that if there is an edge pixel at a given point in space there is likely to be another edge pixel nearby. Creating a new sensor in the neighborhood of a known edge point obviates the need for a remote sensor to take the time to thermally diffuse into the neighborhood. The action of this adaptive edge detection system is shown in figures 5 through 7. Figure 5 shows the initial distribution of sensing elements (with density 0.02). Figure 6 shows the distribution after 10 time steps, and figure 7 the distribution after 100 time steps. It is seen that all of the edges have been detected within 100 time steps.

8 Conclusion

We have introduced a new class of sensing devices, the active point sensor array, which promise an increase in functionality for some applications relative to the current generation of sensing devices. Some of the advantages of active point sensor arrays over conventional sensor arrays are adaptability, efficient use of communication bandwidth, object oriented output, and complex data processing integrated with the sensing elements. Micro-scale analogs to our active point sensor approach exist in natural systems. For example, chemical messaging (e.g. hormones, neurotransmitters, pheromones etc.) demonstrates the use of object oriented, asynchronous, communication, while immunological adaption and exploration illustrates the power of altering the "program" of exploratory active sensors based on their current state. The capabilities of our active point sensor arrays as compared with standard sensor arrays can be thought of in terms of an analogy with the comparison between neural networks and standard computers. The processing power of neural nets is due in large part to the connectivity of their constituents rather than due to the computational complexity of these constituents. Likewise, in active sensor systems the bulk of the computational power comes from the motion of the sensors and not from the sensing elements themselves or from complex operations on the sensory data.

References

[1] Aloimonos, J., Weiss, I., and Bandyopadhyay, "Active vision", Proceedings of the 1987 International Conference on Computer Vision, London, pp 35-54

[2] Bajcsy, R., "Active perception vs. passive perception", Proceedings of the 3rd IEEE Workshop on Computer Vision, Bellaire, pp 55-59, 1985

[3] Blake, A., Zisserman, A., and Knowles, G., "Surface descriptions from stereo and shading", Image and Vision Computing, Vol. 3, No. 4, pp 183-191, 1985 (also published in **Shape From Shading,** Horn, B.K.P. and Brooks, M.J., eds., MIT Press, Cambridge, MA, 1989)

[4] Burnet, F.M., **The Clonal Selection Theory of Immunity,** Vanderbilt University Press, Nashville Tenn., 1959

[5] Burt, P., "Algorithms and architectures for smart sensing", in the proceedings of the 1988 Darpa Image Understanding Workshop, pp 139-153

[6] Clark, J.J. and Ferrier, N.J., "Modal control of an attentive vision system.", Proceedings of the 1988 International Conference on Computer Vision, Tampa Florida, pp 514-523

[7] Clark, J.J., and Friedman, D.J., "Local processing of sensor array data using current mode circuitry", Proceedings of the 1989 Int. Electron Devices Meeting, Washington, D.C., Dec. 1989

[8] Clark, J.J., and Yuille, A.L., **Data Fusion For Sensory Information Processing Systems,** Kluwer Academic Publishers, Boston, 1990

[9] Drexler, K.E., **The Engines of Creation,** Doubleday Anchor Press, New York, 1986

[10] Freeman, H., "Computer processing of line drawing images.", Computer Surveys, Vol. 6, No. 1, 1974, pp 57-98

[11] Friedman, A., **Stochastic Differential Equations and Applications,** Academic Press, New York, 1975

[12] Fujita, H. and Omodaka, A., "Electrostatic actuators for micromechatronics.", Proceedings of the IEEE Workshop on Micro Robots and Teleoperators, Hyannis, MA 1987

[13] Garabedian, P., **Partial Differential Equations,** Wiley, New York, 1964

[14] Gibson, J.J., **The Ecological Approach to Visual Perception,** , Houghton Mifflin, Boston, 1979

[15] Harris, J., Koch, C., and Luo, J., "Resistive fuses: analog hardware for detecting discontinuities in early vision", in **Analog VLSI Implementations of Neural Systems,** Mead, C. and Ismail, M. eds., Kluwer, Norwell, MA 1989

[16] Horn, B.K.P. "Obtaining shape from shading information", Chap. 4, in **The Psychology of Computer Vision,** Ed P.H. Winston. McGraw-Hill Book Company. New York. pp 115-155. 1975.

[17] Kass, M., Witkin, A., and Terzopoulos, D., "Snakes: Active contour models", Proceedings of the 1987 International Conference on Computer Vision, London, pp 259-268

[18] Koch, C., "Seeing chips: analog VLSI circuits for computer vision", Neural Computation, Vol. 1, pp 184-200

[19] Marroquin, J., Mitter, S., and Poggio, T., "Probabilistic Solution of Ill-Posed Problems in Computational Vision", Journal of the American Statistical Association, March 1987, Vol. 82, No. 397, pp. 76-89.

[20] Mead, C.A., **Analog VLSI and Neural Systems,** Addison Wesley, Reading MA, 1989

[21] Perona, P., and Malik, J., "A network for multiscale image segmentation", Proceedings of the 1988 IEEE International Conference on Circuits and Systems, pp 2565-2568

[22] Sivilotti, M., Mahowald, M., and Mead, C., "Real-time visual computation using analog CMOS processing arrays", In 1987 Stanford Conference on VLSI, pp 295-312.

[23] Stansfield, S., "Visually aided tactile exploration", Proceedings of the 1987 IEEE Conference on Robotics and Automation, pp 1487-1492

[24] Trimmer, W.S.N., and Gabriel, K.J., "Design considerations for a practical elctrostatic micromotor.", Sensors and Actuators, Vol. 11, No. 2, pp 189-206, 1987

[25] Yuille, A.L., Cohen, D., and Hallinan, P., "Facial feature extraction by deformable templates", Harvard Robotics Laboratory Technical Report, No. 88-02

[26] Yuille, A.L., Geiger, D., and Bulthoff, H., "Stereo integration, mean field theory and psychophysics", Harvard Robotics Laboratory Technical Report, No. 89-11

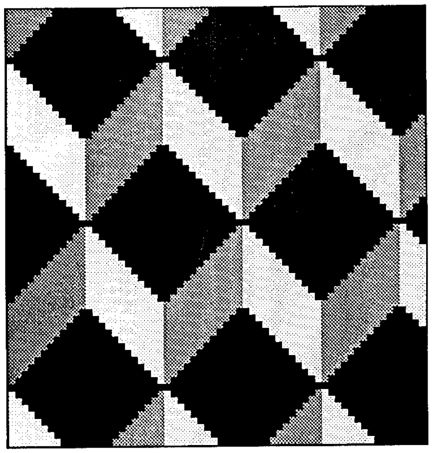

Figure 1. The image field I.

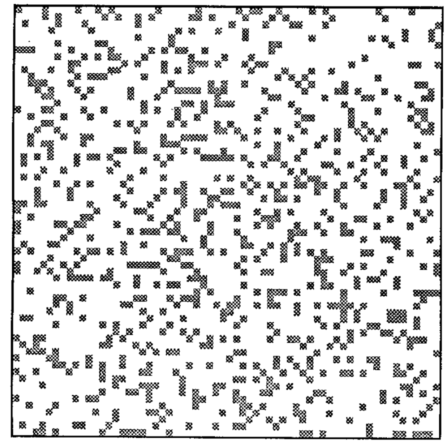

Figure 2. The initial sensor distribution, non-adaptive case.

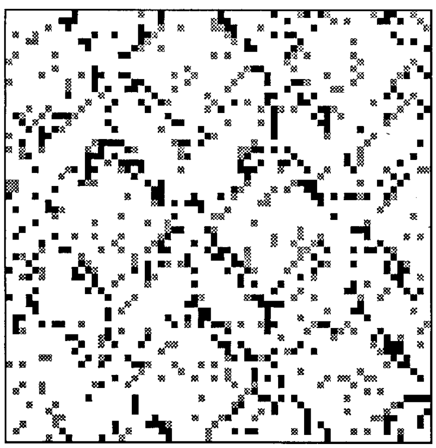

Figure 3. The sensor distribution after 50 time steps.

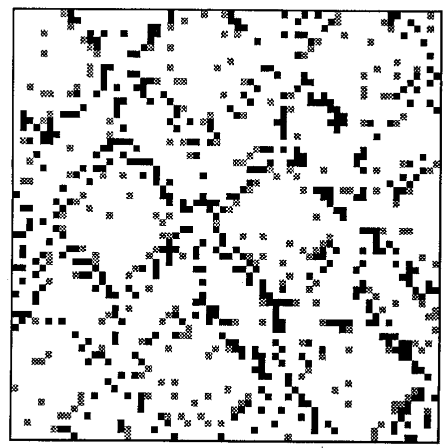

Figure 4. The sensor distribution after 100 time steps.

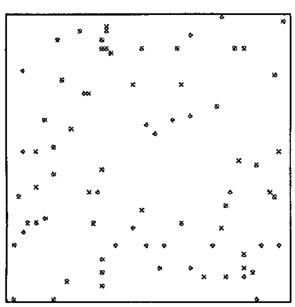

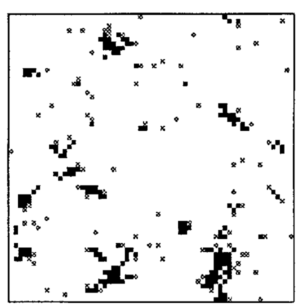

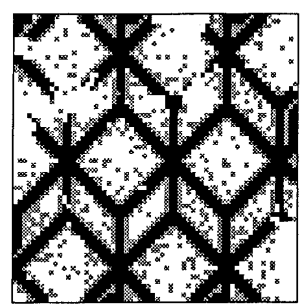

Figure 5. The initial sensor distribution, adaptive case.

Figure 7. The sensor distribution after 100 time steps.

Figure 6. The sensor distribution after 10 time steps.