

Snakes: Active Contour Models

MICHAEL KASS, ANDREW WITKIN, and DEMETRI TERZOPOULOS

Schlumberger Palo Alto Research, 3340 Hillview Ave., Palo Alto, CA 94304

Abstract

A snake is an energy-minimizing spline guided by external constraint forces and influenced by image forces that pull it toward features such as lines and edges. Snakes are active contour models: they lock onto nearby edges, localizing them accurately. Scale-space continuation can be used to enlarge the capture region surrounding a feature. Snakes provide a unified account of a number of visual problems, including detection of edges, lines, and subjective contours; motion tracking; and stereo matching. We have used snakes successfully for interactive interpretation, in which user-imposed constraint forces guide the snake near features of interest.

1 Introduction

In recent computational vision research, low-level tasks such as edge or line detection, stereo matching, and motion tracking have been widely regarded as autonomous bottom-up processes. Marr and Nishihara [11], in a strong statement of this view, say that up to the 2.5D sketch, no “higher-level” information is yet brought to bear: the computations proceed by utilizing only what is available in the image itself. This rigidly sequential approach propagates mistakes made at a low level without opportunity for correction. It therefore imposes stringent demands on the reliability of low-level mechanisms. As a weaker but more attainable goal for low-level processing, we argue that it ought to provide sets of alternative organizations among which higher-level processes may choose, rather than shackling them prematurely with a unique answer.

In this paper we investigate the use of energy minimization as a framework within which to realize this goal. We seek to design energy functions whose local minima comprise the set of alternative solutions available to higher-level processes. The choice among these alternatives could require some type of search or high-level reasoning. In the absence of a well-developed high-level mechanism, however, we use an interactive approach to explore the alternative

organizations. By adding suitable energy terms to the minimization, it is possible for a user to push the model out of a local minimum toward the desired solution. The result is an active model that falls into the desired solution when placed near it.

Energy minimizing models have a rich history in vision going back at least to Sperling's stereo model [16]. Such models have typically been regarded as autonomous, but we have developed interactive techniques for guiding them. Interacting with such models allows us to explore the energy landscape very easily and develop effective energy functions that have few local minima and little dependence on starting points. We hope thereby to make the job of high-level interpretation manageable yet not constrained unnecessarily by irreversible low-level decisions.

The problem domain we address is that of finding salient image contours—edges, lines, and subjective contours—as well as tracking those contours during motion and matching them in stereopsis. Our variational approach to finding image contours differs from the traditional approach of detecting edges and then linking them. In our model, issues such as the connectivity of the contours and the presence of corners affect the energy functional and hence the detailed structure of the locally optimal contour. These issues can, in principle, be resolved by very high-




Fig. 1. Lower-left: Original wood photograph from Brodatz. Others: Three different local minima for the active contour model.

level computations. Perhaps more importantly, high-level mechanisms can interact with the contour model by pushing it toward an appropriate local minimum. Optimization and relaxation have been used previously in edge and line detection, [3,5,13,24,25], but without the interactive guiding used here.

In many image interpretation tasks, the correct interpretation of low-level events can require high-level knowledge. Consider, for example, the three perceptual organizations of two dark lines in figure 1. The three different organizations correspond to three different local minima in our line-contour model. It is important to notice that the shapes of the lines are materially different in the three examples, not just because of a different linking of line segments. The segments themselves are changed by the perceptual organization.

Without detailed knowledge about the object in view, it is difficult to justify a choice among the three interpretations. Knowing that wood is a layered structure, or perhaps inferring its layered structure from elsewhere in the picture could help to rule out interpretation (b). Beyond that, the 'correct' interpretation could be very task dependent. In many domains, such as analyzing seismic data, the choice of interpretation can depend on expert knowledge. Different seismic interpreters can derive significantly different perceptual organizations from the same seismic sections depending on their knowledge and training. Because a single 'correct' interpretation cannot always be defined, we suggest low-level mechanisms which seek appropriate local minima instead of searching for global minima.

Unlike most other techniques for finding salient contours, our model is *active*. It is always

minimizing its energy functional and therefore exhibits dynamic behavior. Because of the way the contours slither while minimizing their energy, we call them snakes. Changes in high-level interpretation can exert forces on a snake as it continues its minimization. Even in the absence of such forces, snakes exhibit hysteresis when exposed to moving stimuli.

Snakes do not try to solve the entire problem of finding salient image contours. They rely on other mechanisms to place them somewhere near the desired contour. However, even in cases where no satisfactory automatic starting mechanism exists, snakes can still be used for semi-automatic image interpretation. If an expert user pushes a snake close to an intended contour, its energy minimization will carry it the rest of the way. The minimization provides a ‘power assist’ for a person pointing to a contour feature.

Snakes are an example of a more general technique of matching a deformable model to an image by means of energy minimization. In spirit and motivation, this idea shares much with the rubber templates of Widrow [23]. From any starting point, the snake deforms itself into conformity with the nearest salient contour. We have applied the same basic techniques to the problem of 3D object reconstruction from silhouettes by using energy minimizing surfaces with preferred symmetries [17]. We expect this general approach will find a wide range of applicability in vision.

In section 2 we present a basic mathematical description of snakes along with their Euler equations. Then in section 3 we give details of the energy terms that can make a snake attracted to different types of important static, monocular features such as lines, edges, and subjective contours. Section 4 addresses the applicability of snake models to stereo correspondence and motion tracking. Finally, section 5 discusses further refinements and directions of our current work.

2 Basic Snake Behavior

Our basic snake model is a *controlled continuity* [18] spline under the influence of image forces and external constraint forces. The internal spline forces serve to impose a piecewise smoothness constraint. The image forces push the snake

toward salient image features like lines, edges, and subjective contours. The external constraint forces are responsible for putting the snake near the desired local minimum. These forces can, for example, come from a user interface, automatic attentional mechanisms, or high-level interpretations.

Representing the position of a snake parametrically by $\mathbf{v}(s) = (x(s), y(s))$, we can write its energy functional as

$$\begin{aligned} E_{\text{snake}}^* &= \int_0^1 E_{\text{snake}}(\mathbf{v}(s)) ds \\ &= \int_0^1 E_{\text{int}}(\mathbf{v}(s)) + E_{\text{image}}(\mathbf{v}(s)) \\ &\quad + E_{\text{con}}(\mathbf{v}(s)) ds \end{aligned} \quad (1)$$

where E_{int} represent the internal energy of the spline due to bending, E_{image} gives rise to the image forces, and E_{con} gives rise to the external constraint forces. In this section, we develop E_{int} and give examples of E_{con} for interactive interpretation. E_{image} is developed in section 3.

2.1 Internal Energy

The internal spline energy can be written

$$E_{\text{int}} = (\alpha(s)|\mathbf{v}_s(s)|^2 + \beta(s)|\mathbf{v}_{ss}(s)|^2)/2 \quad (2)$$

The spline energy is composed of a first-order term controlled by $\alpha(s)$ and a second-order term controlled by $\beta(s)$. The first-order term makes the snake act like a membrane and the second-order term makes it act like a thin plate. Adjusting the weights $\alpha(s)$ and $\beta(s)$ controls the relative importance of the membrane and thin-plate terms. Setting $\beta(s)$ to zero at a point allows the snake to become second-order discontinuous and develop a corner. The controlled continuity spline is a generalization of a Tikonov stabilizer [19] and can formally be regarded as *regularizing* [14,15] the problem.

Details of our minimization procedure are given in the appendix. The procedure is an $O(n)$ iterative technique using sparse matrix methods. Each iteration effectively takes implicit Euler steps with respect to the internal energy and explicit Euler steps with respect to the image and external constraint energy. The numeric con-

siderations are relatively important. In a fully explicit Euler method, it takes $O(n^2)$ iterations each of $O(n)$ time for an impulse to travel down the length of a snake. The resulting snakes are flaccid. In order to erect more rigid snakes, it is vital to use a more stable method that can accommodate the large internal forces. Our semi-implicit method allows forces to travel the entire length of a snake in a single $O(n)$ iteration.

2.2 Snake Pit

In order to experiment with different energy functions for low-level visual tasks, we have developed a user-interface for snakes on a Symbolics Lisp Machine. The interface allows a user to select starting points and exert forces on snakes interactively as they minimize their energy. In addition to its value as a research tool, the user-interface has proven very useful for semiautomatic image interpretation. In order to specify a particular image feature, the user has

only to push a snake near the feature. Once close enough, the energy minimization will pull the snake in the rest of the way. Accurate tracking of contour features can be specified in this way with little more effort than pointing. The snake energy minimization provides a 'power assist' for image interpretation.

Our interface allows the user to connect a spring to any point on a snake. The other end of the spring can be anchored at a fixed position, connected to another point on a snake, or dragged around using the mouse. Creating a spring between $\mathbf{x}_1$ and $\mathbf{x}_2$ simply adds $-k(\mathbf{x}_1 - \mathbf{x}_2)^2$ to the external constraint energy E_{con} .

In addition to springs, the user interface provides a $1/r^2$ repulsion force controllable by the mouse. The $1/r$ energy functional is clipped near $r = 0$ to prevent numerical instability, so the resulting potential is depicted by a volcano icon. The volcano is very useful for pushing a snake out of one local minimum and into another.

Figure 2 shows the snake-pit interface being




Fig. 2. The Snake Pit user-interface. Snakes are shown in black, springs and the volcano are in white.

used. The two dark lines are different snakes which the user has connected with two springs shown in white. The other springs attach points on the snakes to fixed positions on the screen. In the upper right, the volcano can be seen bending a nearby snake. Each of the snakes has a sharp corner which has been specified by the user.

3 Image Forces

In order to make snakes useful for early vision we need energy functionals that attract them to salient features in images. In this section, we present three different energy functionals which attract a snake to lines, edges, and terminations. The total image energy can be expressed as a weighted combination of the three energy functionals

$$E_{\text{image}} = w_{\text{line}}E_{\text{line}} + w_{\text{edge}}E_{\text{edge}} + w_{\text{term}}E_{\text{term}} \quad (3)$$

By adjusting the weights, a wide range of snake behavior can be created.

3.1 Line Functional

The simplest useful image functional is the image intensity itself. If we set

$$E_{\text{line}} = I(x, y) \quad (4)$$

then depending on the sign of w_{line} , the snake will be attracted either to light lines or dark lines. Subject to its other constraints, the snake will try to align itself with the lightest or darkest nearby contour. This energy functional was used with the snakes shown in figure 1. By pushing with the




Fig. 3. Two edge snakes on a pear and potato. Upper-left: The user has pulled one of the snakes away from the edge of the pear. Others: After the user lets go, the snake snaps back to the edge of the pear.

volcano, a user can rapidly move a snake from one of these positions to another. The coarse control necessary to do so suggests that symbolic attentional mechanisms might be able to guide a snake effectively.

3.2 Edge Functional

Finding edges in an image can also be done with a very simple energy functional. If we set $E_{\text{edge}} = -|\nabla I(x, y)|^2$, then the snake is attracted to contours with large image gradients. An example of the use of this functional is shown in figure 3. In the upper left, a user has placed two snakes on the edges of the pear and potato. He has then pulled part of the snake off the pear with a spring. The remaining pictures show what happens when he lets go. The snake snaps back rapidly to the boundary of the pear.

3.3 Scale Space

In figure 3, the snake was attracted to the pear boundary from a fairly large distance away because of the spline energy term. This type of convergence is rather common for snakes. If part of a snake finds a low-energy image feature, the spline term will pull neighboring parts of the snake toward a possible continuation of the feature. This effectively places a large energy well around a good local minimum. A similar effect can be achieved by spatially smoothing the edge- or line-energy functional. One can allow the snake to come to equilibrium on a very blurry energy functional and then slowly reduce the blurring. The result is minimization by scale-continuation [20,21].

In order to show the relationship of scale-space continuation to the Marr-Hildreth theory of edge-detection [10], we have experimented with a slightly different edge functional. The edge-energy functional is

$$E_{\text{line}} = -(G_{\sigma} * \nabla^2 I)^2 \quad (5)$$

where G_{σ} is a Gaussian of standard deviation σ . Minima of this functional lie on zero-crossings of $G_{\sigma} * \nabla^2 I$ which define edges in the Marr-

Hildreth theory. Adding this energy term to a snake means that the snake is attracted to zero-crossings, but still constrained by its own smoothness. Figure 4 shows scale-space continuation applied to this energy functional. The upper left shows the snake in equilibrium at a very coarse scale. Since the edge-energy function is very blurred, the snake does a poor job of localizing the edge, but is attracted to this local minimum from very far away. Slowly reducing the blurring leads the snake to the position shown in the upper right and finally to the position shown in the lower left. For reference, the zero-crossings of $G_{\sigma} * \nabla^2 I$ corresponding to the energy function of the snake in the lower left are shown superimposed on the same snake in the lower right. Note that the snake jumps from one piece of a zero-crossing contour to another. At this scale, the shapes of the zero-crossings are dominated by the small-scale texture rather than the region boundary, but the snake nevertheless is able to use the zero-crossings for localization because of its smoothness constraint.

3.4 Termination Functional

In order to find terminations of line segments and corners, we use the curvature of level lines in




Fig. 4. Upper-left: Edge snake in equilibrium at coarse scale. Upper-right: Snake in equilibrium at intermediate scale. Lower-left: Final snake equilibrium after scale-space continuation. Lower-right: Zero-crossings overlayed on final snake position.

a slightly smoothed image. Let $C(x, y) = G_o(x, y) * I(x, y)$ be a slightly smoothed version of the image. Let $\theta = \tan^{-1}(C_y/C_x)$ be the gradient angle and let $\mathbf{n} = (\cos \theta, \sin \theta)$ and $\mathbf{n}_\perp = (-\sin \theta, \cos \theta)$ be unit vectors along and perpendicular to the gradient direction. Then the curvature of the level contours in $C(x, y)$ can be written

$$E_{\text{term}} = \frac{\partial \theta}{\partial \mathbf{n}_\perp} \quad (6)$$

$$= \frac{\partial^2 C / \partial \mathbf{n}_\perp^2}{\partial C / \partial \mathbf{n}} \quad (7)$$

$$= \frac{C_{yy}C_x^2 - 2C_{xy}C_xC_y + C_{xx}C_y^2}{(C_x^2 + C_y^2)^{3/2}} \quad (8)$$

By combining E_{edge} and E_{term} , we can create a snake that is attracted to edges or terminations. Figure 5 shows an example of such a snake exposed to a standard subjective contour illusion [7]. The shape of the snake contour between the edges and lines in the illusion is entirely determined by the spline smoothness term. The variational problem solved by the snake is very closely related to a variational formulation proposed by Brady et al. [2] for the interpolation of subjective contours. Ullman's [22] proposal of interpolating

using piecewise circular arcs would probably also produce a very similar interpolation. An appealing aspect of the snake model is that the same snake that finds subjective contours can very effectively find more traditional edges in natural imagery. It may, moreover, provide some insight into why the ability to see subjective contours is important.

A further unusual aspect of the snake model that bears on the psychophysics of subjective contours is hysteresis. Since snakes are constantly minimizing their energy, they can exhibit hysteresis when shown moving stimuli. Figure 6 shows a snake tracking a moving subjective contour. As the horizontal line segment on the right moves over, the snake bends more and more until the internal spline forces overpower the image forces. Then the snake falls off the line and reverts to a smoother shape. Bringing the line segment close enough to the snake makes the snake reattach. While it is difficult to show the hysteresis in a still picture, the reader can easily verify the corresponding hysteresis in human vision by recreating the moving stimulus. This type of hysteresis is uncharacteristic of purely bottom-up processes and global optimizations.




Fig. 5. Right: Standard subjective contour illusion. Left: Edge/termination snake in equilibrium on the subjective contour.




Fig. 6. Above left: Dynamic subjective contour illusion. Sequence is left to right, top to bottom. Above Right: Snake attracted to edges and terminations. As the moving horizontal

line slides to the right, the snake bends until it falls off the line. Bringing the line close enough makes the snake reattach.

4 Stereo and Motion

4.1 Stereo

Snakes can also be applied to the problem of stereo matching. In stereo, if two contours correspond, then the disparity should vary slowly along the contour unless the contour rapidly recedes in depth. Psychophysical evidence [4] of a disparity gradient limit in human stereopsis indicates that the human visual system at least to some degree assumes that disparities do not change too rapidly with space. This constraint can be expressed in an additional energy functional for a stereo snake:

$$E_{\text{stereo}} = (\mathbf{v}_s^L(s) - \mathbf{v}_s^R(s))^2 \quad (9)$$

where $\mathbf{v}^L(s)$ and $\mathbf{v}^R(s)$ are left and right snake contours.

Since the disparity smoothness constraint is applied along contours, it shares a strong similarity with Hildreth's [8] smoothness constraint for computing optic flow. This constraint means that during the process of localizing a contour in one eye, information about the corresponding contour in the other eye is used. In stereo snakes, the stereo match actually affects the detection and localization of the features on which the match is based. This differs importantly, for example, from the Marr-Poggio stereo theory [12] in which the basic stereo matching primitive zero-crossings always remain unchanged by the matching process.

Figure 7 shows an example of a 3D surface reconstructed from disparities measured along a




Fig. 7. Bottom: Stereogram of a bent piece of paper. Below: Surface reconstruction from the outline of the paper matched using stereo snakes. The surface model is rendered from a very different viewpoint than the original to emphasize that it is a full 3D model, rather than a 2.5D model.

single stereo snake on the outline of a piece of paper. The surface is rendered from a very different viewpoint than the original to emphasize that a 3D model of the piece of paper has been computed rather than merely a 2.5D model.

4.2 Motion

Once a snake finds a salient visual feature, it 'locks on.' If the feature then begins to move slowly, the snake will simply track the same local minimum. Movement that is too rapid can cause a snake to flip into a different local minimum, but for ordinary speeds and video-rate sampling, snakes do a good job of tracking motion. Figure 8 shows eight selected frames out of a two-second video sequence. Edge-attracted snakes were initialized by hand on the speaker's lips in the first frame. After that, the snakes tracked the lip movements automatically.

The motion tracking was done in this case without any interframe constraints. Introducing such constraints will doubtless make the tracking




Fig. 8. Selected frames from a 2-second video sequence showing snakes used for motion tracking. After being initialized to

the speaker's lips in the first frame, the snakes automatically track the lip movements with high accuracy.

more robust. A simple way to do so is to give the snake mass. Then the snake will predict its next position based on its previous velocity.

5 Conclusion

Snakes have proven useful for interactive specification of image contours. We have begun to use them as a basis for interactively matching 3D models to images. As we develop better energy functionals the 'power assist' of snakes becomes increasingly effective. Scale-space continuation can greatly enlarge the capture region around features of interest.

The snake model provides a unified treatment to a collection of visual problems that have been treated differently in the past. Edges, lines, and subjective contours can all be found by essentially the same mechanisms. Tracking these features through motion and matching them in stereo is easily handled in the same framework.

Snakes, perhaps, embody Marr's notion of 'least commitment' [9] more than his bottom-up 2.5D sketch. The snake provides a number of widely separated local minima to further levels of processing. Instead of committing irrevocably to a single interpretation, snakes can change their interpretation based on additional evidence from higher levels of processing. They can, for example, adjust monocular edge-finding based on binocular matches.

We believe that the ability to have all levels of visual processing influence the lowest-level visual interpretations will turn out to be very important. Local energy-minimizing systems like snakes offer an attractive method for doing this. The energy minimization leaves a much simpler problem for higher level processing.

Acknowledgments

Kurt Fleischer helped greatly with the snake pit user-interface and created the all-important volcano icon. John Platt helped us develop snake theory and guided us around the infinite abyss of numerical methods. Snake-lips Atkinson provided the visual motion stimulus.

Appendix: Numerical Methods

Let $E_{\text{ext}} = E_{\text{image}} + E_{\text{con}}$. When $\alpha(s) = \alpha$, and $\beta(s) = \beta$ are constants, minimizing the energy functional of equation (1) gives rise to the following two independent Euler equations:

$$\alpha x_{ss} + \beta x_{ssss} + \frac{\partial E_{\text{ext}}}{\partial x} = 0 \quad (10)$$

$$\alpha y_{ss} + \beta y_{ssss} + \frac{\partial E_{\text{ext}}}{\partial y} = 0 \quad (11)$$

When $\alpha(s)$ and $\beta(s)$ are not constant, it is simpler to go directly to a discrete formulation of the energy functional in equation (2). Then we can write

$$E_{\text{snake}}^* = \sum_{i=1}^n E_{\text{int}}(i) + E_{\text{ext}}(i) \quad (12)$$

Approximating the derivatives with finite differences and converting to vector notation with $\mathbf{v}_i = (x_i, y_i) = (x(ih), y(ih))$, we expand $E_{\text{int}}(i)$

$$E_{\text{int}}(i) = \alpha_i |\mathbf{v}_i - \mathbf{v}_{i-1}|^2 / 2h^2 + \beta_i |\mathbf{v}_{i-1} - 2\mathbf{v}_i + \mathbf{v}_{i+1}|^2 / 2h^4 \quad (13)$$

where we define $\mathbf{v}(0) = \mathbf{v}(n)$. Let $f_x(i) = \partial E_{\text{ext}} / \partial x_i$ and $f_y(i) = \partial E_{\text{ext}} / \partial y_i$ where the derivatives are approximated by a finite difference if they cannot be computed analytically. Now the corresponding Euler equations are

$$\begin{aligned} \alpha_i (\mathbf{v}_i - \mathbf{v}_{i-1}) - \alpha_{i+1} (\mathbf{v}_{i+1} - \mathbf{v}_i) \\ + \beta_{i-1} [\mathbf{v}_{i-2} - 2\mathbf{v}_{i-1} + \mathbf{v}_i] \\ - 2\beta_i [\mathbf{v}_{i-1} - 2\mathbf{v}_i + \mathbf{v}_{i+1}] \\ + \beta_{i+1} [\mathbf{v}_i - 2\mathbf{v}_{i+1} + \mathbf{v}_{i+2}] \\ + (f_x(i), f_y(i)) = 0 \end{aligned} \quad (14)$$

The above Euler equations can be written in matrix form as

$$\mathbf{A}\mathbf{x} + \mathbf{f}_x(\mathbf{x}, \mathbf{y}) = 0 \quad (15)$$

$$\mathbf{A}\mathbf{y} + \mathbf{f}_y(\mathbf{x}, \mathbf{y}) = 0 \quad (16)$$

where $\mathbf{A}$ is a pentadiagonal banded matrix.

To solve equations (15) and (16), we set the right-hand sides of the equations equal to the product of a step size and the negative time derivatives of the left-hand sides. Taking into ac-

count derivatives of the external forces we use requires changing $\mathbf{A}$ at each iteration, so we achieve faster iteration by simply assuming that $\mathbf{f}_x$ and $\mathbf{f}_y$ are constant during a time step. This yields an explicit Euler method with respect to the external forces. The internal forces, however, are completely specified by the banded matrix, so we can evaluate the time derivative at time t rather than time $t - 1$ and therefore arrive at an implicit Euler step for the internal forces. The resulting equations are

$$\mathbf{A}\mathbf{x}_t + \mathbf{f}_x(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) = -\gamma(\mathbf{x}_t - \mathbf{x}_{t-1}) \quad (17)$$

$$\mathbf{A}\mathbf{y}_t + \mathbf{f}_y(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) = -\gamma(\mathbf{y}_t - \mathbf{y}_{t-1}) \quad (18)$$

where γ is a step size. At equilibrium, the time derivative vanishes and we end up with a solution of equations (15) and (16).

Equations (17) and (18) can be solved by matrix inversion:

$$\mathbf{x}_t = (\mathbf{A} + \gamma\mathbf{I})^{-1}(\mathbf{x}_{t-1} - \mathbf{f}_x(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})) \quad (19)$$

$$\mathbf{y}_t = (\mathbf{A} + \gamma\mathbf{I})^{-1}(\mathbf{y}_{t-1} - \mathbf{f}_y(\mathbf{x}_{t-1}, \mathbf{y}_{t-1})) \quad (20)$$

The matrix $\mathbf{A} + \gamma\mathbf{I}$ is a pentadiagonal banded matrix, so its inverse can be calculated by *LU* decompositions in $O(n)$ time [6,1]. Hence equations (19) and (20) provide a rapid solution to equations (15) and (16). The method is implicit with respect to the internal forces, so it can solve very rigid snakes with large step sizes. If the external forces become large, however, the explicit Euler steps of the external forces will require much smaller step sizes.

References

1. A. Benson, and D.J. Evans, ACM TRANS. MATHEMATICAL SOFTWARE, vol. 3, pp. 96-103, 1977.
2. M. Brady, W.E.L. Grimson, and D. Langridge, "Shape encoding and subjective contours," PROC. AM. ASSOC. ARTIF. INTEL., Stanford University, 1980.
3. D.J. Burr, "Elastic matching of line drawings," IEEE TRANS. PAMI-8, p. 708, 1986.
4. P. Burt, and B. Julesz, "A disparity gradient limit for binocular fusion," SCIENCE, vol. 208, pp. 615-617, 1980.
5. M.A. Fischler, and R.A. Elschlager, "The representation and matching of pictorial structure," IEEE TRANS. ON COMPUTERS, vol. C-22, pp. 67-92, 1973.
6. I. Gladwell, and R. Wait (eds.), A SURVEY OF NUMERICAL METHODS FOR PARTIAL DIFFERENTIAL EQUATIONS. Clarendon: Oxford, 1979.
7. Kanisza, "Subjective contours," SCIENTIFIC AMERICAN, vol. 234, pp. 48-52, 1976.
8. E. Hildreth, "The computation of the velocity field," PROC. ROY. SOC. (LONDON), vol. B221, pp. 189-220.
9. D. Marr, VISION. Freeman: San Francisco, 1982.
10. D. Marr and E. Hildreth, "A theory of edge detection," PROC. ROY. SOC. (LONDON), vol. B207, pp. 187-217, 1980.
11. D. Marr and H.K. Nishihara, "Visual information processing: Artificial Intelligence and the sensorium of sight," TECHNOLOGY REVIEW, vol. 81(1), October 1978.
12. D. Marr and T. Poggio, "A computational theory of human stereo vision," PROC. ROY. SOC. (LONDON), vol. B204, pp. 301-328, 1979.
13. A. Martelli, "An application of heuristic search methods to edge and contour detection," CACM, vol. 19, p. 73, 1976.
14. T. Poggio and V. Toree, "Ill-posed problems and regularization analysis in early vision," PROC. AARPA Image Understanding Workshop, New Orleans, L.A., Baumann (ed.), pp. 257-263, 1984.
15. T. Poggio, V. Toree, and C. Koch, "Computational vision and regularization theory," NATURE, vol. 317(6035), pp. 314-319, 1985.
16. G. Sperling, "Binocular vision: a physical and a neural theory," AM. J. PSYCHOLOGY, vol. 83, pp. 461-534, 1970.
17. D. Terzopoulos, A. Witkin, and M. Kass, "Symmetry-seeking models for 3D object reconstruction," INT. J. COMPUTER VISION, vol. 3, 1987.
18. D. Terzopoulos, "Regularization of inverse visual problems involving discontinuities," IEEE TRANS. PAMI-8, p. 413, 1986.
19. A.N. Tikhonov, "Regularization of incorrectly posed problems," SOV. MATH. DOKL., vol. 4, pp. 1624-1627, 1963.
20. A. Witkin, "Scale space filtering," PROC. EIGHTH INT. JOINT CONF. ARTIF. INTEL., Karlsruhe, pp. 1019-1021, 1983.
21. A. Witkin, D. Terzopoulos, and M. Kass, "Signal matching through scale space," PROC. AM. ASSOC. ARTIF. INTEL., Philadelphia, pp. 714-719, 1986.
22. S. Ullman, "Filling the gaps: The shape of subjective contours and a model for their generation," BIOLOGICAL CYBERNETICS, vol. 25, 1976.
23. B. Widrow, "The rubber mask technique, parts I and II," PATTERN RECOGNITION, vol. 5, pp. 175-211, 1973.
24. S. Zucker, R. Hummel, and A. Rosenfeld, "An application of relaxation labeling to line and curve enhancement," IEEE TRANS. ON COMPUTERS, vol. C-26, p. 394, 1977.
25. S. Zucker, "Computational and psychophysical experiments in grouping: Early orientation selection." In HUMAN AND MACHINE VISION, Jacob Beck, et al. (eds.), Academic Press: New York, pp. 545-567, 1983.