

Model Driven Edge Detection

P. Fua and Y. G. Leclerc*

SRI International

333 Ravenswood Avenue, Menlo Park, CA 94025

(fua@ai.sri.com leclerc@ai.sri.com)

Machine Vision and Applications, **3**, 1990

Abstract

Standard edge detectors fail to find most relevant edges, finding either too many or too few, because they lack a geometric model to guide their search. We present a technique that integrates both photometric and geometric models with an initial estimate of the boundary. The strength of this approach lies in the ability of the geometric model to overcome various photometric anomalies, thereby finding boundaries that could not otherwise be found. Furthermore, edges can be scored based on their goodness of fit to the model, thus allowing one to use semantic model information to accept or reject the edges.

1 Introduction

In real-world images, object boundaries cannot be detected solely on the basis of their photometry because of the presence of noise and various photometric anomalies. Thus, all methods for finding boundaries based on purely local statistical criteria are bound to err, finding either too many or too few edges based on arbitrary thresholds.

To supplement the weak and noisy local information, we consider the geometric constraints that object models can provide. We do this by describing a boundary as an elastic curve with a deformation energy derived from the geometric constraints, as suggested previously (Terzopoulos 1987; Kass et al. 1988). Local minima in this energy correspond to boundaries that match the object model exactly. We incorporate the photometric constraints by defining photometric energy as the average along the curve of a scalar energy field derived from the photometric model of an edge. Local minima in this energy correspond to boundaries that match the photometric model exactly. We find a candidate boundary by deforming the curve in such a way as to minimize its total energy, which is the sum of the deformation and photometric energies. We call this deformation process “optimizing” the curve. Once a curve has been optimized, i.e., once it has settled in a local minimum of the total energy, thus effecting a compromise between the two constraints. We can then utilize the object models to determine if the curve actually corresponds to an object boundary.

Such “energy minimizing curves” have two key advantages:

- The geometric constraints are directly used to guide the search for a boundary.

*The work reported here was partially supported by the Defense Advanced Research Projects Agency under contracts DACA76-85-C-0004 and MDA903-86-C-0084. We wish to thank Demetri Terzopoulos for the many discussions we have had concerning this work.

- The edge information is integrated along the entire length of the curve, providing a large support *without including the irrelevant information of points not on the curve*.

Taken together, they allow energy minimizing curves to find photometrically weak boundaries that local edge detectors simply could not find without also finding many irrelevant boundaries.

Energy minimizing curves, as we call them, were originated by Terzopoulos, Kass, and Witkin as “snakes” (Terzopoulos 1987; Kass et al. 1988). This paper was motivated by the need to solve two major problems with snakes. First, the precise relationship between the local minima found by these snakes and the standard definitions of an edge had not been elucidated. Second, snakes were basically designed for use in interactive environments. Consequently, they require fairly extensive parameter adjustments from one image to the next, and even from one part of an image to the next.

In this paper, we offer three major contributions. (1) We show the precise relationship between the local minima of the energy minimizing curve and the standard definitions of an edge. (2) We present a procedure for automatically adjusting the parameters of these curves. (3) We briefly show how the relationship and the procedure can be combined in an automated system that extracts roads and buildings from aerial images, described extensively in (Fua and Hanson 1988).

2 A Photometric Model: Step Edges

Several authors (Haralick 1984; Canny 1986; Torre and Poggio 1986) define step-edge points as those for which the first directional derivative of image intensity, in the direction of the gradient, is extremal. That is, an edge point (x_0, y_0) satisfies

$$\left. \frac{\partial^2 \mathcal{I}(x, y)}{\partial \mathbf{g}_0^2} \right|_{(x, y) = (x_0, y_0)} = 0, \quad (1)$$

where

$$\begin{aligned} \frac{\partial}{\partial \mathbf{g}_0} &= \mathbf{g}_0 \cdot \nabla, \\ \mathbf{g}_0 &= \mathbf{g}(x_0, y_0) = \frac{\nabla \mathcal{I}(x_0, y_0)}{|\nabla \mathcal{I}(x_0, y_0)|}, \end{aligned}$$

and $\mathcal{I}(x, y)$ represents the image intensities after convolution with a twice differentiable kernel (typically a Gaussian).¹ It can be shown that Equation (1) is equivalent to

$$\left. \frac{\partial |\nabla \mathcal{I}(x, y)|}{\partial \mathbf{g}_0} \right|_{(x, y) = (x_0, y_0)} = 0.$$

In other words, the Haralick *et al.* criterion for an ideal step-edge point is equivalent to the criterion that the magnitude of the image intensity gradient be maximal in the gradient direction.

¹This definition is a good approximation to the position of discontinuities in the underlying intensity function when the diameter of support of the operator is much smaller than the radius of curvature of the edge. In particular, it is not applicable at corners and junctions of edges.

However, for noisy non-ideal step edges the gradient direction may become very unreliable, especially at some distance from the edge itself. We therefore use the following, somewhat more general definition of an edge:

Definition *An edge is a curve \mathcal{C} whose points have a gradient magnitude that is maximal in the direction normal to the curve. That is, all points along the curve (called edge points) satisfy*

$$\frac{\partial |\nabla \mathcal{I}(\mathbf{f}(s))|}{\partial \mathbf{n}(\mathbf{f}(s))} = 0, \quad (2)$$

where s is the arc-length of \mathcal{C} , $\mathbf{f}(s)$ is a vector function mapping the arc-length s to points (x, y) in the image, and $\mathbf{n}(\mathbf{f}(s))$ is the normal to \mathcal{C} .

Note that this definition is equivalent to the Haralick *et al.* definition for ideal noise-free step edges.

To incorporate this photometric model in the energy minimizing curve approach, we define the photometric energy of the curve \mathcal{C} as

$$\mathcal{E}_P(\mathcal{C}) = -\frac{1}{|\mathcal{C}|} \int_0^{|\mathcal{C}|} |\nabla \mathcal{I}(\mathbf{f}(s))| ds$$

where $|\mathcal{C}|$ is the length of \mathcal{C} . For simplicity of notation, the limits on the above integral will not be repeated in the remainder of the text.

To understand the relationship between optimized curves and edges, consider the following. When \mathcal{C} is either an open or closed curve that has been optimized, i.e., when it is a local minimum of \mathcal{E}_P with respect to infinitesimal deformations of the curve, we prove in the Appendix that:

$$\frac{\partial |\nabla \mathcal{I}(\mathbf{f}(s))|}{\partial \mathbf{n}(\mathbf{f}(s))} = \gamma(s) \left(|\nabla \mathcal{I}(\mathbf{f}(s))| - \frac{1}{|\mathcal{C}|} \int |\nabla \mathcal{I}(\mathbf{f}(s))| ds \right), \quad (3)$$

where $\gamma(s)$ is the curvature of \mathcal{C} . Furthermore, we prove that when \mathcal{C} is an open curve that has been optimized,

$$|\nabla \mathcal{I}(\mathbf{f}(0))| = |\nabla \mathcal{I}(\mathbf{f}(|\mathcal{C}|))| = \frac{1}{|\mathcal{C}|} \int |\nabla \mathcal{I}(\mathbf{f}(s))| ds \quad (4)$$

also holds. These two equations have the following consequences.

Equation (3) implies that the points along an optimized curve are all edge points when either its curvature is zero (i.e., it is a straight line segment), or the magnitude of the gradient along the curve is constant. *Consequently, optimized curves whose curvature is small or for which the magnitude of the gradient is approximately constant are good approximations to edges.* In particular, we have found that replacing $|\nabla \mathcal{I}(x, y)|$ by $\log |\nabla \mathcal{I}(x, y)|$ (which has no effect for ideal step edges) yields more stable solutions because the second term in Equation (3) is smaller for non-constant edge strengths, and hence the approximation is generally better.

For an open-ended curve, Equation (4) must also be satisfied. This equation implies that the curve is stable only when the gradient magnitude at the end points equals the average. In other words, an open ended energy minimizing curve placed exactly along an edge will shrink or expand, following the edge, until this condition is met. When the gradient magnitude is constant, of course, this condition is always met, and curves therefore remain unchanged. Thus, again, using $\log |\nabla \mathcal{I}(x, y)|$ yields more stable solutions because this equation is more nearly satisfied for non-constant edge strengths. An important note is

that if we had chosen to use the integral of gradient magnitude, as opposed to the average, this equation would not hold: curves would simply expand without bound.

After optimization, candidate edges can be scored on how well they match the photometric model by computing the proportion of points along the curve that satisfy Equation (2) to within some tolerance. We use this score in our applications to accept or reject any given edge.

3 A Geometric Model: Smooth Curves

3.1 Theory

We have seen in the previous section that energy minimizing curves match edges well whenever the curves have a low curvature. To ensure stable results we define the deformation energy of such curves so that their curvature remains small and the minimum energy state is close to the photometric edge. This can be achieved simply by defining the deformation energy of a curve \mathcal{C} as:

$$\mathcal{E}_D(\mathcal{C}) = \int \gamma(s)^2 ds$$

where $\gamma(s)$ is the curvature of \mathcal{C} . The total energy of such a curve is then $\mathcal{E}(\mathcal{C}) = \mathcal{E}_P(\mathcal{C}) + \lambda \mathcal{E}_D(\mathcal{C})$, where $\mathcal{E}_P(\mathcal{C})$ is the photometric energy defined in the previous section.

To delineate a photometric edge using one of these curves, one must provide an initial estimate of the location of the curve and then optimize it using the procedure described in the next subsection. Previously (Kass et al. 1988), λ was chosen interactively. We have found that a single value of λ , however carefully chosen, is inadequate for use across a wide variety of images. Instead, we compute λ in the following adaptive manner.

Let \mathcal{C}_i be the initial curve and \mathcal{C}_f be the curve after optimization. Let δ_i denote the operator

$$\delta_i = \left. \frac{\delta}{\delta \mathcal{C}} \right|_{\mathcal{C}=\mathcal{C}_i},$$

that is, δ_i is the variational operator evaluated at the initial curve. Define δ_f similarly for the final curve. Since, by definition, $\delta_f \mathcal{E}(\mathcal{C}) = 0$, then

$$\lambda = \frac{|\delta_f \mathcal{E}_P(\mathcal{C})|}{|\delta_f \mathcal{E}_D(\mathcal{C})|}.$$

When the initial estimate is known to be close to the final answer, we choose

$$\lambda = \frac{|\delta_i \mathcal{E}_P(\mathcal{C})|}{|\delta_i \mathcal{E}_D(\mathcal{C})|}$$

so that we remain as close to the initial estimate as possible. If instead we want to smooth the initial estimate, we can use higher values of λ . This is in fact the simplest way of imposing a geometric model, in this case smoothness, upon the data.

3.2 Implementation

In the actual implementation, the curves are described as polygons with n equidistant vertices $X = \{(x_i, y_i), i = 1, \dots, n\}$ and the deformation energy can be made discrete as

$$\mathcal{E}_D = \sum_i (2x_i - x_{i-1} - x_{i+1})^2 + (2y_i - y_{i-1} - y_{i+1})^2 \quad (5)$$

when $\gamma(s)$ is small. To perform the optimization we could use a simple gradient descent technique (Leclerc and Fua 1987), but it would be extremely slow for curves with a large number of vertices. Instead, we embed the curve in a viscous medium and solve the equation of dynamics

$$\frac{\partial \mathcal{E}}{\partial X} + \alpha \frac{dX}{dt} = 0,$$

where $\mathcal{E} = \mathcal{E}_P + \lambda \mathcal{E}_D$ and α is the viscosity of the medium (Terzopoulos 1987). Since the deformation energy \mathcal{E}_D in Equation 5 is quadratic, its derivative with respect to X is linear and therefore

$$\frac{\partial \mathcal{E}_D}{\partial X} = K X,$$

where K is a pentadiagonal matrix. Thus, each iteration of the optimization amounts to solving the linear equation:

$$K X_t + \alpha(X_t - X_{t-1}) = \left. \frac{\partial \mathcal{E}_P}{\partial X} \right|_{X_{t-1}}. \quad (6)$$

Because K is pentadiagonal, the solution to this set of equations can be computed efficiently in $O(n)$ time using LU decomposition and backsubstitution. Note that the LU decomposition need be recomputed only when α changes. When α is constant, only the backsubstitution step is required. For a curve with 100 vertices, the LU decomposition step takes .05 seconds on a 3675 Symbolics Lisp Machine with a Floating Point Accelerator, and the backsubstitution step takes only .012 seconds.

Previously (Terzopoulos 1987), α was chosen interactively, and remained fixed throughout the optimization process. Since this has proven to be inadequate for use across a variety of images, we use the following adaptive procedure for computing α .

We start with an initial step size Δ and compute the viscosity so that

$$\alpha = \frac{\sqrt{2n}}{\Delta} \left| \frac{\partial \mathcal{E}}{\partial X} \right|,$$

where n is the number of vertices. This ensures that the initial displacement of each vertex is on the average of magnitude Δ . Because of the non-linear term, we must verify that the energy has decreased from one iteration to the next. If, instead, the energy has increased, the curve is reset to its previous position, the step size is decreased, and the viscosity recomputed accordingly. This is repeated until the step size becomes less than some threshold value. In most cases, because of the presence of the linear term that propagates constraints along the whole curve in one iteration, it only takes a small number of iterations to optimize the initial curve. For example, going from the initial estimate of the closed curve with about 20 vertices shown in Figure 1a to the optimized result shown in Figure 1b took 20 iterations and 3.9 seconds on 3675 Symbolics Lisp Machine with a Floating Point Accelerator . Only 18 iterations and 2.8 seconds were required for the open curve. Note that the time actually spent solving the equations of the dynamics represents only a relatively small fraction of the total time required to perform the complete optimization; most of the time is actually spent updating the snakes and fetching the image values.

Figure 1

3.3 Evaluation Criterion

To illustrate the utility of our edge-point criterion (Equation 2), consider the synthetic image shown in Figure 2, wherein one boundary of the central region is smooth while the

Figure 2

Figure 3

other is jagged. Two energy minimizing curves were entered interactively (Figure 3b) and optimized using the smoothness constraint above (Figure 3c). Removing all points along the optimized curves that do not satisfy Equation 2 yields Figure 3d.

While most of the pixels belonging to the smooth boundary satisfy the edge criterion, very few of those belonging to the jagged boundary do. Also note that a small part of the topmost curve failed to settle correctly along the edge because of noise. However, this failure is easily identified because it is the only part of the curve for which a significant fraction of the points fails to satisfy the criterion.

Thus, computing the proportion of edge points in the optimized edges is an efficient and reliable way to measure how well the photometric boundaries match the smoothness constraint.

4 Applications

In this section we present two applications of our energy-minimizing curves. First we describe an application of our model to road delineation in an interactive context, where the user provides a rough initial estimate of the location of the contours. This interactive procedure has been integrated in the SRI Cartographic Modeling Environment (Hanson and Quam 1988). It allows a cartographer to roughly sketch features of interest, and then interactively adjust the resultant optimized contours. We then outline an application of our model to road and building delineation in the context of the fully automated system described in details in (Fua and Hanson 1988).

4.1 Interactive Road Delineation

We can apply the smooth curve model of the previous section directly to the problem of finding road boundaries in aerial images. We model roads as ribbons whose smoothly curved edges are approximately parallel. A ribbon is implemented as a polygonal curve forming the center of the road. Associated with each vertex i of this curve is a width w_i that defines the two curves that are the candidate road boundaries. The photometric energy is defined as the sum of the photometric energies of the two boundary curves, and the deformation energy is the sum of the smoothness term described above and an additional term that enforces the parallel constraint:

$$\sum_i (w_i - w_{i-1})^2.$$

Figure 4

We have applied this model to two road segments in the aerial image of Figure 4a. Figure 4b shows the initial estimates for two road segment boundaries and Figure 4c shows the final optimized boundaries.

This image is an excellent example of a situation where a local edge detector such as the Canny edge detector is insufficient, by itself, to find the relevant boundaries. Because the roads are dirt roads with ill-defined edges, no single edge-strength threshold yields a satisfactory set of boundaries: when the threshold is too high, most of the road edges are lost, whereas when the threshold is too low, there is a plethora of irrelevant edges (see Figure 5).

Figure 5

In the next section, we briefly describe an automated system that exploits the power of energy minimizing curves to overcome this problem.

4.2 Automated Road and Building Delineation

In our automated system (Fua and Hanson 1988), roads are modeled as above, while buildings are modeled as polygons whose edges meet at right angles.

We begin by using the synthetic image of Figure 2 to demonstrate how the system uses energy minimizing curves to extract smooth and straight edges from a set of Canny edge maps. Since a full description of the system is beyond the scope of this paper, we can only briefly outline the subsequent steps that lead to the generation of road and building outlines, and show the results of this system applied to an aerial image.

Figure 6

Smooth Edges: The system begins by applying the Canny edge detector to the image using several thresholds (Figure 6). The pixels in each of these maps are then linked using a linker (Fischler and Wolf 1983). Each of these linked segments is used as the starting point for an energy minimizing curve. After optimization, the curves are broken into segments wherever sharp bends are found. Those segments for which at least 70% of the points satisfy the edge criterion are accepted as candidates for the next step (Figure 7a,b,c). This 70% figure has been used for all examples in this paper, and indeed for every other application of the automated system. Experimentally, curves with less than 50% of the points satisfying the edge criterion are almost always spurious, while curves with more than 90% are almost always correct. We chose 70% as a compromise that rejects most of the spurious edges without losing a significant number of correct ones.

Figure 7

Since the same edge may be discovered, in whole or in part, in several of the Canny edge maps, the system must choose an appropriate subset of these curves. The basic goal of this selection process is to choose the longest possible curves that do not overlap and have good edge quality. This goal is achieved by selecting the largest subset of non-overlapping curves that maximize the number of pixels satisfying the edge criterion minus the number of pixels that do not. The result of this selection process is shown in Figure 7d.

This technique provides an effective way of selecting the smooth boundaries in the image. Furthermore by combining the output of the Canny edge detector using several arbitrary thresholds, we have eliminated the necessity of carefully choosing an appropriate threshold. This is a substantial advantage since, as discussed previously, choosing this threshold is in general a very difficult task.

Figure 8

Straight Edges: The system begins with the same linked Canny edge images as before, finds points of high curvature in the linked segments, and fits straight lines between those points. The straight lines are used as the starting points for energy minimizing curves with only two vertices. As previously, the curves are optimized and only those with 70% of points satisfying the edge criterion are retained (Figure 8a,b,c). A subset of these curves is selected using the same procedure as above (Figure 8d). Note that, in this particular case, all of the final curves could have been extracted from a single Canny edge map. This rarely occurs in practice because the noise across a real image is never as uniform as it is in this synthetic image.

Figure 9

Figure 10

Automatic Road Delineation: Applying the smooth edge technique to the aerial image of Figure 9 yields the edges shown in Figure 10a. These edges are then grouped into sets of parallel edges, each of which is used as the starting point for a ribbon, as described in Section 4.1. After optimization, the system uses a selection criterion to retain only the best

road candidates (Figure 10c). The groups from which these candidates arose are shown in Figure 10b.

Figure 11

Automatic Building Delineation: Applying the straight edge technique to the same aerial image yields the edges shown in Figure 11a. These edges are then grouped into rectilinear edge structures, and closed using a simple edge tracker. Finally, the outline of these rectilinear contours is optimized by using an energy minimizing curve and a rectilinearity constraint. Once again, the system uses a selection criterion to retain only the best building candidates (Figure 11c). The groups from which these candidates arose are shown in Figure 11b.

5 Summary and Conclusion

We have presented a technique for finding object boundaries that integrates both photometric and geometric models with an initial estimate of the boundary. The models are incorporated by defining an energy function that is minimal when the models are satisfied exactly. The initial estimate is used as the starting point for finding a local minimum of this energy function, by embedding the initial curve in a viscous medium and solving the equations of dynamics.

The strengths of this “energy minimizing curve” approach are that the geometric constraints are directly used to guide the search for a boundary, and that the edge information is integrated along the entire length of the curve, thereby providing a large support.

We have shown the precise relationship between optimized curves and a standard definition of an edge, and that this relationship can be used to determine when an optimized curve should be accepted as a candidate edge for further processing. We have also presented methods for automatically choosing the optimization parameters and how they should change over time. We found that these contributions were crucial to our ability to embed energy minimizing curves in the automated system briefly discussed here and presented in detail elsewhere (Fua and Hanson 1988); these methods have provided us with the ability to find photometrically weak boundaries that local edge detectors simply could not find without also finding many irrelevant boundaries.

Appendix

In this appendix, we show that a necessary and sufficient condition for an open curve \mathcal{C} to be a local extremum of

$$\mathcal{E}(\mathcal{C}) = \frac{\int G(\mathbf{f}(s))ds}{\int ds} \quad (1)$$

with respect to all infinitesimal deformations is that

$$\frac{dG(\mathbf{f}(s))}{d\mathbf{n}(s)} = \gamma(s) \left[G(\mathbf{f}(s)) - \frac{\int G(\mathbf{f}(s))ds}{\int ds} \right] \quad (2)$$

and

$$G(\mathbf{f}(0)) = G(\mathbf{f}(|\mathcal{C}|)) = \frac{\int G(\mathbf{f}(s))ds}{\int ds}. \quad (3)$$

where the curve \mathcal{C} is parameterized by its arc-length s ; $\mathbf{f}(s)$ is a vector function mapping the arc-length s to points (x, y) along the curve; $\mathbf{n}(s)$ is the normal to the curve; $\mathbf{t}(s)$ is the

tangent to the curve; and $\gamma(s)$ is its curvature. Furthermore, eq. (2) alone is the necessary and sufficient condition for closed curves.

To prove this result, consider deformations of the curve \mathcal{C} , which we shall call \mathcal{C}_λ , such that the mapping from arc-length s to points (x, y) is of the form

$$\mathbf{f}_\lambda(s) = \mathbf{f}(s) + \lambda (\eta(s)\mathbf{n}(s) + \tau(s)\mathbf{t}(s)), \quad (4)$$

where $\eta(s)$ and $\tau(s)$ are arbitrary continuous and differentiable function. When \mathcal{C} is a closed curve, $\eta(s)$ and $\tau(s)$ are constrained to be periodic of period $|\mathcal{C}|$. Since \mathcal{C} is a local extremum of $\mathcal{E}(\mathcal{C})$ if and only if

$$\left. \frac{d\mathcal{E}(\mathcal{C}_\lambda)}{d\lambda} \right|_{\lambda=0} = 0, \quad (5)$$

for all $\eta(s)$ and $\tau(s)$, we now prove the following theorem.

Theorem: Equation (5) holds for all $\eta(s)$ and $\tau(s)$ if and only if eq. (2) and (3) above are satisfied.

Proof: By definition,

$$\mathcal{E}(\mathcal{C}_\lambda) = \frac{\int G(\mathbf{f}_\lambda(s_\lambda)) ds_\lambda}{\int ds_\lambda}, \quad (6)$$

where s_λ is the arc-length of \mathcal{C}_λ . We can rewrite this in terms of s by using the equality

$$ds_\lambda = |\mathbf{f}'_\lambda(s)| ds. \quad (7)$$

Thus,

$$\mathcal{E}(\mathcal{C}_\lambda) = \frac{\int G(\mathbf{f}_\lambda(s)) |\mathbf{f}'_\lambda(s)| ds}{\int |\mathbf{f}'_\lambda(s)| ds}, \quad (8)$$

and therefore, to within the non-zero factor of $1/(\int |\mathbf{f}'_\lambda(s)| ds)^2$:

$$\begin{aligned} \frac{d\mathcal{E}(\mathcal{C}_\lambda)}{d\lambda} &= \frac{d}{d\lambda} \left[\int G(\mathbf{f}_\lambda(s)) |\mathbf{f}'_\lambda(s)| ds \right] \int |\mathbf{f}'_\lambda(s)| ds \\ &\quad - \left[\int G(\mathbf{f}_\lambda(s)) |\mathbf{f}'_\lambda(s)| ds \right] \frac{d}{d\lambda} \int |\mathbf{f}'_\lambda(s)| ds \\ &= \left[\int \frac{dG(\mathbf{f}_\lambda(s))}{d\lambda} |\mathbf{f}'_\lambda(s)| + G(\mathbf{f}_\lambda(s)) \frac{d|\mathbf{f}'_\lambda(s)|}{d\lambda} ds \right] \int |\mathbf{f}'_\lambda(s)| ds \\ &\quad - \left[\int G(\mathbf{f}_\lambda(s)) |\mathbf{f}'_\lambda(s)| ds \right] \int \frac{d|\mathbf{f}'_\lambda(s)|}{d\lambda} ds. \end{aligned} \quad (9)$$

The first term we need to evaluate in eq. (9) is $d|\mathbf{f}'_\lambda(s)|/d\lambda$, for which we need the following equalities:

$$\begin{aligned} \frac{d\mathbf{f}(s)}{ds} &= \mathbf{f}'(s) = \mathbf{t}(s) \\ \frac{d\mathbf{t}(s)}{ds} &= \mathbf{t}'(s) = \gamma(s)\mathbf{n}(s) \\ \frac{d\mathbf{n}(s)}{ds} &= \mathbf{n}'(s) = -\gamma(s)\mathbf{t}(s). \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbf{f}'_\lambda(s) &= \mathbf{f}'(s) + \lambda [\eta'(s)\mathbf{n}(s) + \eta(s)\mathbf{n}'(s) + \tau'(s)\mathbf{t}(s) + \tau(s)\mathbf{t}'(s)] \\ &= \mathbf{t}(s) + \lambda [\eta'(s)\mathbf{n}(s) - \eta(s)\gamma(s)\mathbf{t}(s) + \tau'(s)\mathbf{t}(s) + \tau(s)\gamma(s)\mathbf{n}(s)] \\ &= \mathbf{t}(s) [1 + \lambda(\tau'(s) - \eta(s)\gamma(s))] + \mathbf{n}(s) [\lambda(\eta'(s) + \tau(s)\gamma(s))]. \end{aligned} \quad (10)$$

Since $\mathbf{t}(s)$ and $\mathbf{n}(s)$ are, by definition, orthogonal unit vectors, we have

$$|\mathbf{f}'_\lambda(s)| = \sqrt{[1 + \lambda(\tau'(s) - \eta(s)\gamma(s))]^2 + [\lambda(\eta'(s) + \tau(s)\gamma(s))]^2}, \quad (11)$$

and therefore

$$\frac{d|\mathbf{f}'_\lambda(s)|}{d\lambda} = \frac{[1 + \lambda(\tau'(s) - \eta(s)\gamma(s))](\tau'(s) - \eta(s)\gamma(s)) + \lambda(\eta'(s) + \tau(s)\gamma(s))^2}{\sqrt{[1 + \lambda(\tau'(s) - \eta(s)\gamma(s))]^2 + [\lambda(\eta'(s) + \tau(s)\gamma(s))]^2}}. \quad (12)$$

The second term we need to evaluate is

$$\begin{aligned} \frac{dG(\mathbf{f}_\lambda(s))}{d\lambda} &= (\eta(s)\mathbf{n}(s) + \tau(s)\mathbf{t}(s)) \cdot \nabla G(\mathbf{f}_\lambda(s)) \\ &= \eta(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{n}(s)} + \tau(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{t}(s)}. \end{aligned} \quad (13)$$

Substituting (11), (12), and (13) into (9), evaluating at $\lambda = 0$, and multiplying by the non-zero factor $1/(\int |\mathbf{f}'_\lambda(s)| ds)^2$:

$$\begin{aligned} \frac{d\mathcal{E}(\mathcal{C}_\lambda)}{d\lambda} \Big|_{\lambda=0} = 0 &= \int \eta(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{n}(s)} + \tau(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{t}(s)} + G(\mathbf{f}(s))(\tau'(s) - \eta(s)\gamma(s)) ds \int ds \\ &\quad - \int G(\mathbf{f}(s)) ds \int (\tau'(s) - \eta(s)\gamma(s)) ds. \end{aligned}$$

Dividing by $\int ds$ and rearranging terms, we have

$$\begin{aligned} 0 &= \int \eta(s) \left[\frac{dG(\mathbf{f}(s))}{d\mathbf{n}(s)} - \gamma(s) \left(G(\mathbf{f}(s)) - \frac{\int G(\mathbf{f}(s)) ds}{\int ds} \right) \right] ds \\ &\quad + \int \tau(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{t}(s)} + \tau'(s) \left[G(\mathbf{f}(s)) - \frac{\int G(\mathbf{f}(s)) ds}{\int ds} \right] ds. \end{aligned}$$

A necessary and sufficient condition for the first integral to be zero for all $\eta(s)$ is that the term multiplying $\eta(s)$ be zero, which is eq. (2).

Integrating the second integral by parts:

$$\int \tau(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{t}(s)} ds + \tau(s) \left(G(\mathbf{f}(s)) - \frac{\int G(\mathbf{f}(s)) ds}{\int ds} \right) \Big|_0^{|\mathcal{C}|} - \int \tau(s) \frac{dG(\mathbf{f}(s))}{d\mathbf{t}(s)} ds = 0, \quad (14)$$

therefore,

$$\tau(|\mathcal{C}|) \left(G(\mathbf{f}(|\mathcal{C}|)) - \frac{\int G(\mathbf{f}(s)) ds}{\int ds} \right) - \tau(0) \left(G(\mathbf{f}(0)) - \frac{\int G(\mathbf{f}(s)) ds}{\int ds} \right) = 0. \quad (15)$$

A necessary and sufficient condition for this equality to hold for all $\tau(s)$ is that both terms be zero, which is eq. (3).

Thus, a necessary and sufficient condition for an open curve \mathcal{C} to be a local extremum of $\mathcal{E}(\mathcal{C})$ is that both equations (2) and (3) hold.

Furthermore, for closed curves, $\tau(0) \equiv \tau(|\mathcal{C}|)$ and $G(\mathbf{f}(0)) \equiv G(\mathbf{f}(|\mathcal{C}|))$, hence eq. (15) is always satisfied. Therefore, a necessary and sufficient condition for a closed curve \mathcal{C} to be a local extremum of $\mathcal{E}(\mathcal{C})$ is that eq. (2) holds.

Q.E.D.

References

- [1] Canny J (1986) A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Vision*, (8)(6):679–698
- [2] Fischler M, Wolf H (1983) Linear delineation. *Proceedings of the IEEE Conference on Computer Vision and Image Processing*, Washington, D.C., 351–356.
- [3] Fua P, Hanson A (1988) Extracting Generic Shapes Using Model-Driven Optimization. *Proceedings of the 1988 DARPA Image Understanding Workshop*, Boston, MA, 994–1004
- [4] Haralick RM (1984) Digital step edges from zero crossings of second directional derivatives. *IEEE Transactions on Pattern Analysis and Machine Vision*, (6)(1):58–68
- [5] Hanson AJ and Quam L(1988), “Overview of the SRI Cartographic Modeling Environment,” in *Proceedings of the Image Understanding Workshop*, Boston, MA, 576–582.
- [6] Kass M, Witkin A, Terzopoulos D (1988) Snakes: active contour models. *International Journal of Computer Vision*, (1)(4):321–331
- [7] Leclerc YG, Fua P (1987) Finding object boundaries using guided gradient ascent. *Topical Meeting on Machine Vision*, Technical Digest Series, Optical Society of America, Washington, D.C., (12):168–161
- [8] Terzopoulos D (1987) On matching deformable models to images. *Topical Meeting on Machine Vision*, Technical Digest Series, Optical Society of America, Washington, D.C., (12):160–167
- [9] Torre V and Poggio T (1986) On edge detection. *IEEE Transactions on Pattern Analysis and Machine Vision*, (8)(2):147–163