

Active Contours using Finite Elements to Control Local Scale

Persephoni Karaolani,
G D Sullivan and K D Baker
Intelligent Systems Group,
Department of Computer Science
University of Reading, RG6 2AY, UK.
p.karaolani@reading.ac.uk.

Abstract

Finite elements allow smoothness to be enforced on the measurement of the image-dependent term in active contours. This improves the stability of the solution, with less computational cost than is incurred by increasing the number of elements. Performance is best when the size of the element matches the scale of the image detail sought.

This property of Finite elements can be used deliberately to select the scale of the image detail recovered by the active contour. The method offers a way to reduce the sensitivity of active contours to localised noise.

1 Introduction.

Active contours (“snakes” and “bubbles”) provide global solutions to the problem of finding continuous edges in images [1, 2, 3, 4, 5, 10]. Instead of building connected edges from point evidence, obtained by purely local operators, the active contour has a pre-established structure (stiff, elastic strings in the case of “snakes”; elastic loops, under internal pressure forces, in the case of “bubbles”). Starting from an initial state, the contour moves under the influence of local forces derived from the image, until these are in equilibrium with the contour’s internal structural forces.

Equilibria occur at minima of the energy potential of the forces acting on the contour. The energy comprises internal and external factors. In typical applications of active contours in vision, the internal forces represent the contour’s elasticity and stiffness; the external force is provided by the negative of the magnitude of the gradient image, acting along the contour. The former forces tend to make the contour short and smooth; the latter deflects it towards regions of the image with high gradient (i.e. edges). The energy term to be minimised can be expressed as:

$$E(v(s)) = \int_0^1 \left(\alpha \left| \frac{dv}{ds} \right|^2 + \beta \left| \frac{d^2v}{ds^2} \right|^2 + \gamma P_{im}(v) \right) ds \quad (1)$$

where: $v(s) = (x(s), y(s))$ defines a parametric curve in the image plane.

The three terms represent the elastic energy, the stiffness energy, and the potential due to external forces, which in our models is represented by the image intensity function. The parameters α , β , and γ are weights provided to balance the relative importance of the terms in (1).




Figure 1: "Bubbles" solved by Finite Differences

1.1 Finite differences solution

The original method to minimise the energy in Equation (1) was developed using the Euler method with Finite Differences (FD) [1, 2]. This represents the contour as a string of control points, connected to each other by an elastic, stiff link. The image forces act at the control points, where the local derivative of the potential (itself usually a derivative - to emphasise edges) provides a force, deflecting the contour from its natural shape. Each control point is therefore subject to a force dependent on the image at the control point, together with forces due to its 2 immediate neighbours (representing elasticity) and its nearest 4 neighbours (representing stiffness). The global solution to the set of local problems is given by solving two sets of n_d equations in n_d unknowns (the x & y coordinates of the n_d control points). Each equation involves 5 unknowns, which leads to a penta-diagonal system of equations. The problem is profoundly non-linear, in that the image forces are strictly local, and change in an uncontrolled way as the contour moves. The system of equations is therefore repeated iteratively, each iteration moving a little way towards the preferred solution, until stability results.

There are many difficulties with this approach which have limited its application in practice. Firstly, the image forces are only sampled at the control points; therefore in order to make the contour sensitive to fine detail in the image, n_d must be large. Secondly, with large n_d , the gap between control points is small, so that the effect of stiffness (spreading only over 5 points) is only felt on a local scale; this means that a contour is easily caught by local details in an image, which may not be part of the desired global structure - e.g. random noise.

These defects of the Finite Differences method are illustrated in Figure 1, using the implementation reported in [10]. Figure 1(a) shows an original image of a mouth; our purpose is to find and outline the outer edge of the lips. Figure 1(b) shows the negative of the magnitude of the gradient image: we seek to identify the irregular ridge of dark points around the lips. An elliptical "bubble" was initially located just inside the line of

the lips, and the balance of elastic and pressure forces caused it to expand, until the control points were caught on strong local minima (Figures 1(c-f)).

In Figure 1(c), there are 32 control points (shown as +) which have found good individual minima, but the bubble ignores the image lying between them. In Figure 1(e) there are 1024 control points (not shown) and the solution follows the local edges more accurately, but the contour easily becomes snagged on local detail. The problem of snagging can be overcome in part by increasing the weight of the stiffness term (Figure 1(d & f)), but then fine detail (e.g. at the corner of the mouth) becomes even poorer.

The problem is the familiar one of spatial scale. The image force is proportional to a second order derivative of the image (change in the gradient image), which is inherently noisy. Noise can only be combated successfully by integrating information using a scale of smoothness appropriate to the structure being sought; this is highly context-dependent, e.g. the corners of the mouth require finer scale than the lips. As normally implemented, the stiffness and elasticity coefficients in Equation (1) apply globally (though in our implementation a simple method of biasing the equilibrium shape away from symmetry is available [10]). The finite differences method gives no natural way to impose a sense of scale locally.

1.2 Finite element method

An alternative approach to solving the non-linear energy minimisation problem relies on the use of the Finite Element Method (FEM) [9]. Here, the domain s of the contour $v(s)$ is divided into a number of elements, in which the elastic, stiffness and image forces seek their own local minima. Within each element the position of the contour $(x(s), y(s))$ is described by a low order polynomial in s ; successive elements are connected together by constraints which seek to enforce low-order continuity in x & y at their links. A popular choice of function is given by Hermite Cubics, third order curves having fixed end-points and fixed tangents at their ends, which can therefore be connected to each other with first order continuity, while retaining second order continuity within each element [6]. This approach leads to the problem of solving a system of $2n_e$ simultaneous equations in $2n_e$ unknowns, where n_e is the number of elements. An important difference between the FEM and FD approaches is that the image forces act along the entire element, not merely at discrete control points.

The global minimum is sought (as with FD) by moving in the direction of the solution of a system of linear equations determined by the current position of the contour, and iterating. At each iteration, the force acting on each element is recomputed, according to the new position in the image of the element. As before, the process is iterated until the contour is stable.

A key issue in the FEM is the method used to integrate the image forces along an element. A simple scheme ([5]) merely approximated the integral by sampling the image along the element at a fixed number of points. However, this method runs into the same problems encountered by FD: to be sensitive to fine scale in the image, we need many small elements, and in coarser scale parts of the image these easily become tangled by irrelevant local image detail.

More accurate ways to integrate the image forces numerically are available. A widely used scheme is the Gauss-Legendre n -point rule for choosing the sample points for numerical integration, and their relative weights [7]. The order of the rule refers to the number of sample points. If this is high then the estimate of the integral takes account of more image detail, so that noise has a smaller effect on the overall shape of the element (which is still constrained to the cubic form).




Figure 2: Examples of FEM bubbles applied to images having coarse (top) and fine (bottom) scales. Left: Initial position; Middle: 12 elements; Right: 48 elements.

The order can be varied according to the length of the elements: in a simple scheme described here, we make it proportional to the length of the element computed in pixel units, thus (approximately) ensuring regular steps along the contour of approximately 1 pixel distance. This provides a natural way to guarantee that the image forces pay good attention to all the detail along the locally-cubic description of the contour. We also weight the contribution of each element in proportion to its length; this helps to ensure that elements do not form clusters. By deliberately varying the length of the elements it therefore becomes possible to control the scale of smoothness adopted by the contour - and this can be done on a local basis.

2 FEM bubble with adaptable integration.

2.1 Outline of method

The function $v(s)$ is represented as a Hermite Cubic polynomial over each element. The nodes between elements are shown as + in the relevant Figures. The image energy term in Equation (1) is derived by numerical integration along the Hermite Cubic, using the Gauss-Legendre rule [7]. In each element an approximation to the length of the element (in the (x,y) space) is computed, and the number of sample points is made approximately proportional to the length in pixels. The gaussian rule determines the sample points (in the s space) for optimal sampling. The derivative of the image is then computed at these points, and integrated over s . It is known that the Gauss-Legendre rule is optimal for approximating integrals since among all integration rules requiring n_g integrand evaluations, it has the highest possible order of accuracy $(2n_g-1)$ [7].




Figure 3: Cpu times [SPARC2] for one iteration of the energy minimisation calculation, using different number of elements (n_e) and gauss-points (n_g). The dashed lines link cases with equal numbers of integration points ($n_g * n_e$)

At each iteration, the method leads to a hepta-diagonal system of equations in $2n_e$ unknowns, which is currently solved by a conjugate gradient algorithm, though linear methods would be faster (e.g. Cholesky [8]). A scaling factor is introduced to weight elements in proportion to their lengths in the iterated problem, which prevents elements which collapse from having excessive influence.

2.2 Imposing local smoothness

Figure 2 illustrates bubbles consisting of 12 and 48 elements, applied to synthetic images of a blurred modulated circle with added gaussian noise (peak signal is 72, and the noise sd was 20 intensity levels). In each case, the bubble was initially positioned as shown in Figure 2(a & b) for the 12 element version. The images illustrate two different scales of spatial detail. Best performance is obtained (Figure 2(c & f)) when the scale of the element matches the scale of the image. If elements are too coarse (Figure 2(d)), then the bubble cannot follow the image; if elements are too fine, (Figure 2(e)) then the system of equations is unnecessarily large, and the elements become over-sensitive to local noise (note the irregular spacing, as elements have collapsed onto noise points).

2.3 Computational costs

Integration along the elements is less costly than increasing the number of elements. Figure 3 shows the results of a simulated experiment in which the computational cost per iteration (averaged over 1000 iterations) was measured as a function of the number of elements and the number of gauss points used to compute the integral. The algorithm is close to linear on the number of elements. However, the number of image points used in the computation ($= n_g * n_e$) - and hence the opportunity to integrate out image noise - can be doubled with significantly less cpu time cost, by increasing the number of gauss points (moving vertically to the next dashed line in Figure 3).

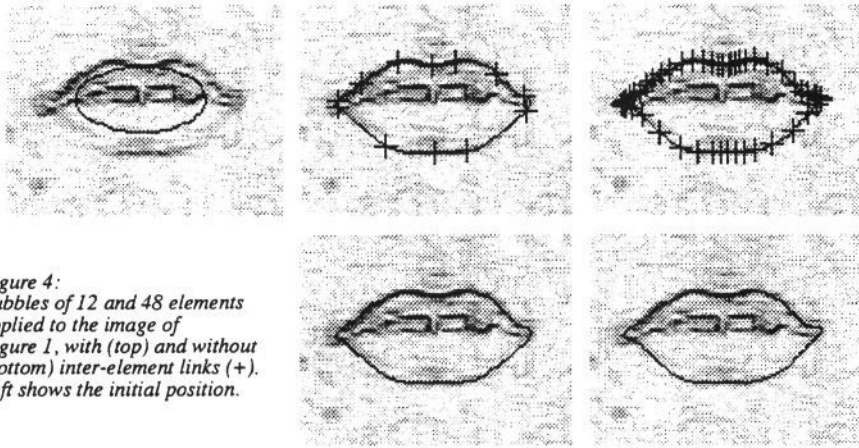


Figure 4:
 Bubbles of 12 and 48 elements
 applied to the image of
 Figure 1, with (top) and without
 (bottom) inter-element links (+).
 Left shows the initial position.

2.4 Demonstration on real images

Figure 4 shows the 12 and 48 element bubbles applied to the magnitude of the gradient image of lips in Figure 1(b), from the initial starting position shown. In comparison with the FD method of Figure 1, the FEM seems more accurate; it is also far less sensitive to initial conditions. The 48 element bubble has more accurate performance at high curvature points (e.g. the corners of the mouth), but the 12 element bubble recovers a smoother description of the slowly varying regions (edges of lips).

3 Conclusions.

We have demonstrated a new approach to active contours, using FEM with an adaptable order of integration. This provides a mechanism to integrate the local forces accurately. A major benefit is that fuller use is made of the local information, and smoothness can be enforced, without becoming excessively sensitive to noise; a secondary benefit is that the scale of the detail sought can be imposed on a local basis.

The scheme is also efficient. The well-known FD approach [1, 2] has a computational cost that increases linearly with the number of control points (those points where the image is sampled). If the size of an element is well-matched to the local spatial scale, then many fewer elements are needed than FD control points. The information loss is avoided by interpolating the element, within the constraint of the cubic form; this represents a smaller computational burden, since finer sampling does not increase the number of unknowns in the linearised problem. In iterative schemes, computing speed is critical. We have found that the FEM requires far fewer iterations than the FD method for convergence, and in our implementations (the FD on an 1024 processor SIMD multiprocessor, the FEM on a solitary SUN), the FEM is also far faster, and produces better results from simple initial conditions, though direct comparisons on more similar implementations have not been made.

The effect of the scaling of the forces by the length of the element is to prevent very small elements (perhaps contracting onto a purely local feature) from contributing as much to the global solution as longer elements. This encourages elements to preserve their lengths, and to avoid elements collapsing onto isolated noise points. In




Figure 5: Multi-scale FEM with 30 elements applied to a synthetic image:
 Left: starting position; Right: final position.

consequence, it is possible to prescribe the initial scale of the contour locally. Figure 5 illustrates performance on a synthetic multi-scale image.

In practical uses of bubbles for image analysis, it will often be the case that some knowledge is available of the expected structure of the objects of interest. The two greatest determinants of the form adopted by an active contour are its preferred (blank image) state, and its initial conditions. The adaptable integration scheme provides a convenient way to create active models which mimic the expected smoothness of images of objects as well as their initial expected positions and shapes.

4 References

- [1] Kass, M., Witkin, A., and Terzopoulos D. Snakes: Active Contour Models, *Int. J. Computer Vision* 1. 321 - 331, 1987
- [2] Terzopoulos D., Witkin A., Kass M. Constraints on Deformable Models: Recovering 3D shape and nonrigid Motion. *AI*, 36, 91-123, 1988
- [3] Amini, A.A., Tehrani S., and Weymouth T.E. Using dynamic programming for minimising the energy of active contours in the presence of hard constraints, In *Second International Conference on Computer Vision* p 95-99, Computer Society Press, December 1988
- [4] Karaolani, P., Sullivan, G. D., Baker, K.D., Baines, M. J. A Finite Element Method for Deformable Models, *Proceedings of the fifth Alvey Vision Conference*, 1989
- [5] Karaolani, P., Sullivan, G. D., Baker, K.D. Parabolic and Hermite Cubic Elements: A Flexible Technique for Deformable Models, *Proceedings of the British Machine Vision Conference*, 325-330, Oxford, September 1990
- [6] Strang, G., and Fix, G. J. *An Analysis of the Finite Element Method*, Prentice-Hall series in automatic computation, N.J. 1973
- [7] Press, W. H., Flannery, B. P., Teukolsky, S. A., Vetterling W. T. *Numerical Recipes*, Cambridge University Press, 1986
- [8] Golub, G. H. and Van Loan C.F. *Matrix Computations*, Hopkins Press, Baltimore, 1990
- [9] Zienkiewicz, C.O., and Morgan, K. *Finite Elements and approximation*, John Wiley & Sons, 1983
- [10] Sullivan, G D, Worrall, A D, Hockney, R W, and Baker K D. Active Contours in Medical Image Processing using a Networked SIMD Array Processor. *British Machine Vision Conference*, 1990, pp 395-400.