

# Active Contours for Head Boundary Extraction by Global and Local Energy Minimisation

Steve R. Gunn and Mark S. Nixon

Image, Speech and Intelligent Systems Group  
Department of Electronics and Computer Science  
University of Southampton

9 November 1997

## Abstract

*Active contours are an attractive choice to extract the head boundary, for deployment within a face recognition or model-based coding scenario. However, conventional snake approaches can suffer difficulty in initialisation and parameterisation. A dual active contour configuration using dynamic programming has been developed to resolve these difficulties by using a global energy minimisation technique and a simplified parameterisation, to enable a global solution to be obtained. The merits of conventional gradient descent based snake (local) approaches, and search based (global) approaches are discussed. In application to find head and face boundaries in front-view face images, the new technique employing dynamic programming is deployed to extract the inner face boundary, along with a conventional normal-driven contour to extract the outer (head) boundary. The extracted contours appear to offer sufficient discriminatory capability for inclusion within an automatic face recognition system.*

**Keywords:** *Active Contours, Head Boundary Extraction, Energy Minimisation, Dynamic Programming*

## 1 Introduction

Active contours are an attractive choice for extraction of boundaries in images. However, the standard difficulties associated with conventional formulations of active contours concern initialisation and parameterisation. The final extracted contour is very sensitive to the position of the initial contour, and the choice of the controlling parameters. A dual active contour has been developed (Gunn and Nixon, 1997), to relieve these difficulties. This reduces problems with contour initialisation by requiring the target feature to lie within an outer and an inner contour, and reduces the number of controlling parameters to a single regularisation parameter, controlling the compromise between the snake's properties and the selected image functional, and a termination parameter. The technique uses two contours which are both formulated using gradient descent minimisation together with an adaptive driving force, so as to seek a global energy minimum between them. A weakness of the approach is that whilst it offers improved performance compared with standard active contour techniques, it cannot guarantee a global minimum solution. In this paper we retain the concept of a dual contour for initialisation but use dynamic programming, to guarantee a global energy minimum. Dynamic programming

is an optimal exhaustive minimisation technique, unlike other global minimisation strategies such as simulated annealing (Rueckert and Burger, 1995; Storvik, 1994). Amongst previous active contour methods implemented with dynamic programming, one implemented the original gradient descent minimisation to find local energy minima (Amini et al., 1990), whereas others have aimed for a global solution (Geiger et al., 1995; Lai and Chin, 1995). Dynamic programming is particularly suited to the minimisation of the snake since dependency is local within the snake model. In this paper we focus on the merits of the conventional snake approaches and the search-based approaches. The paper commences with a description of the new dual search-based minimisation approach in contrast to gradient descent snake minimisation. A comparison between the two techniques is discussed with emphasis on discretisation, the contour space and convergence. The techniques are then applied to extract the face boundary from front view images. Results of application to a face database are described which highlight the potential for use in automatic face recognition.

## 2 Snakes

Active contour techniques can be divided into two groups: gradient descent (local) and search-based (global) approaches. Gradient descent approaches use an iterative technique to refine an initial solution, whereas search-based techniques search for a global minimum.

### 2.1 Gradient Descent Approach (Local)

The original snake model was introduced by Kass et al. (1988). A contour is described parametrically by  $\mathbf{v}(s) = (x(s), y(s))$  where  $x(s)$ ,  $y(s)$  are  $x$ ,  $y$  co-ordinates along the contour and  $s \in [0, 1]$  is normalised arc length. The snake model defines the energy of a contour  $v(s)$ , the snake energy  $E_{snake}$ , to be

$$E_{snake}(\mathbf{v}(s)) = \int_{s=0}^1 \lambda E_{int}(\mathbf{v}(s)) + (1 - \lambda) E_{image}(\mathbf{v}(s)) ds \quad (1)$$

where  $E_{int}$  is the internal energy of the contour, imposing continuity and curvature constraints,  $E_{image}$  is the image energy constructed to attract the snake to desired feature points in the image, and  $\lambda \in [0, 1]$  is the regularisation parameter governing the compromise between adherence to the internal forces and the external image data. The image functional,

$$E_{image} = - \frac{|\nabla I(x, y)| - \min_{x,y} |\nabla I(x, y)|}{\max_{x,y} |\nabla I(x, y)| - \min_{x,y} |\nabla I(x, y)|}, \quad (2)$$

attracts the snakes to edges, and is normalised to reduce the influence of image contrast on the choice of  $\lambda$ . An initial contour evolves by minimising Equation 1 using a gradient descent technique. A characteristic of the original snake model (Kass et al., 1988) was that if a snake was not submitted to any image forces it would contract to a point. To enhance versatility, Cohen and Cohen (1993) proposed an additional normal force which could be applied to the contour. Recent work by Xu et al. (1994) and Gunn and Nixon (1997) has developed alternative schemes for unifying expanding and contracting contours, by removing the internal contraction force without affecting the regularising property. As well as the regularisation parameter, gradient descent techniques using a continuous space require specification of a termination parameter which halts the descent when a chosen level of accuracy is reached.

## 2.2 Search Approach (Global)

A weakness of the gradient descent, or local minimum, approach is the sensitivity to initialisation and difficulty in determining suitable parameters. This can be exaggerated by noise. To overcome this problem we propose a technique which searches for a global minimum within a specified region, constraining the initialisation as in our earlier dual active contour (Gunn and Nixon, 1997). This region is described by two initial contours which define a search space for the technique. Dynamic programming is used to search the space for the optimum solution, which is illustrated in Figure 1. In contrast with the original dual contour a termination parameter is no longer required since the




FIGURE 1: Dynamic Programming Contour Space

technique searches a finite space for a global minimum. A discrete contour is defined by,  $\mathbf{v}_i = (x_i, y_i)$  for  $i = 0 \dots N - 1$  where subscripted arithmetic is modulo  $N$  for closed contours. Each contour point is constrained to lie on a line joining the two initial contours, Figure 1. Each line is discretised into  $M$  points for the dynamic programming search. An open contour is especially compatible with minimisation via dynamic programming. The energy of an open snake is given by,

$$E_{snake}(\mathbf{v}) = \sum_{i=0}^{N-3} E_i(\mathbf{v}_i, \mathbf{v}_{i+1}, \mathbf{v}_{i+2}), \quad (3)$$

emphasising the local dependencies, since a point is dependent only on its immediate neighbours for its energy. The energy of a closed snake contains two extra energy terms arising from the joining of the open snake and the limits of Equation 3 is adjusted to  $N - 1$  accordingly. The energy at each snake point  $\mathbf{v}_i$  is given by

$$E_{i-1}(\mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i+1}) = \lambda E_{int}(\mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i+1}) + (1 - \lambda) E_{image}(\mathbf{v}_i), \quad (4)$$

where  $\lambda \in [0, 1]$  is the regularisation parameter, which controls the compromise between fidelity to the image data and adherence to the snake model. This is the only parameter for which a value needs to be chosen. The number of points for the search space,  $M$ , is determined by the spacing between the two contours with respect to the image co-ordinates. The internal energy function is given by

$$E_{int}(\mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i+1}) = \frac{|\mathbf{v}_{i-1} - 2\mathbf{v}_i + \mathbf{v}_{i+1}|^2}{|\mathbf{v}_{i-1} - \mathbf{v}_{i+1}|^2} \quad (5)$$

The numerator is the discrete curvature term from a conventional snake. The continuity term is less important because point spacing is controlled by constraining the points to lie on the specified lines. The denominator ensures that the internal energy is scale invariant giving no prejudice to large or small contours, only smooth ones. In order to apply dynamic programming to Equation 4, a two element vector of state variables,  $(\mathbf{v}_{i+1}, \mathbf{v}_i)$ , is calculated at each stage. The optimal value function,  $S_i$ , is a function of two adjacent points on the contour and is calculated as

$$S_i(\mathbf{v}_{i+1}, \mathbf{v}_i) = \min_{\mathbf{v}_{i-1}} S_{i-1}(\mathbf{v}_{i+1}, \mathbf{v}_i) + \lambda E_{int}(\mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i+1}) + (1 - \lambda) E_{ext}(\mathbf{v}_i), \quad (6)$$

given the initial conditions  $S_0(\mathbf{v}_1, \mathbf{v}_0) = 0$ . In addition to the energy matrix corresponding to the optimal value function, a position matrix is also required. Each entry of the position matrix at stage  $i$  stores the position of  $\mathbf{v}_{i-1}$  that minimises Equation 6. The optimality function, Equation 6, is evaluated for  $i = 1 \dots N - 2$ . The result is obtained by back-tracking through the position matrix.

The efficiency of dynamic programming is compromised when applied to closed contour problems. To guarantee a global minimum, using the method of (Geiger et al., 1995), requires a separate optimisation to be calculated for all values of  $v_0$  and  $v_1$ , incurring an  $M^2$  increase in complexity over the open contour optimisation. To avoid this increase we propose an approximate solution using a two stage technique which transforms the problem into two open contour optimisations. The method employs a heuristic which states: *removing the closure constraint from the end point rarely effects the mid contour solution*. As illustrated in Figure 2, first an open contour solution is found, which does not apply any continuity or smoothness constraints at the ends. The two points at the mid point of this contour are then taken as the start and end points for the closed contour. A second optimisation of the energy function given in Equation 3 is computed with the fixed  $v_0$  and  $v_1$ . The optimality function, Equation 6, is evaluated for  $i = 1 \dots N$ . By fixing the two points  $\mathbf{v}_0$  and  $\mathbf{v}_1$  the closed contour optimisation can be achieved. Accordingly only two open contour optimisations are required, as opposed to  $M^2$  for the method of (Geiger et al., 1995).




FIGURE 2: Dynamic Programming for Closed Contours

### 2.3 Discretisation

Both snake approaches are usually implemented by discretising the contour into  $N$  points. This is based on an approximation of the integral of Equation 1 by a summation over suitably small segments. The approximation of the contour determines the fidelity of the extracted contours. This discretisation

affects both the internal energy and the image energy. For the approximation to be valid the contour points should have an even spacing. In the gradient descent approach, employing a continuous contour space, even spacing is included by the internal energy. This is only a soft constraint and in regions of high curvature and strong image energy large deviations can occur. In the search based approach the contour spacing is enforced by the discrete contour space, and hence care must be taken in the construction of the search space to provide approximately even contour point spacing. The image energy is usually sampled at the discrete contour points. If the number of points is too small, the approximation to the integral of Equation 1 will be poor and the contour image energy will be a bad approximation. An alternative approach employs continuous contours which inherently provide a complete description of the contour, either by using finite elements (Cohen and Cohen, 1993), local polygonal approximations (Rueckert and Burger, 1995), or B-splines (Menet et al., 1990). These allow the image energy to be integrated along the length of the contour, rather than at discrete points, and render the image energy insensitive to the discretisation of the contour. However this approach is computationally expensive and a similar, but faster, approach can be approximated with a discrete contour, by low-pass filtering the image. This effectively measures the average energy in a local neighbourhood of the contour point.

## 2.4 Convergence

The main advantage of a discrete contour space is that the whole search space is finite and consequently minimisation techniques can exploit this property, providing guaranteed convergence. Gradient descent techniques are not suited to a discrete based grid because of problems with the introduction of extra local minima within the internal energy function. A discrete space can cause a contour to stop as a consequence of an imposed grid, when a continuous model would continue to contract. The gradient descent approach in a continuous contour space stops when the image forces are in equilibrium with the internal forces and any additional driving forces. Accordingly the internal and driving forces can displace the contour solution from a minimum of the image energy. The solution is independent of the depth of the minimum but dependent upon the gradient of the sides of the minimum. A search-based technique does not require a driving force, or the calculation of the gradient of the snake energy function, and therefore solutions are not displaced. In conclusion, search based techniques are compatible with a discrete space and gradient descent techniques are compatible with a continuous space.

## 2.5 Comparison of Approaches

Example simulation results are shown in Figure 3. These concern extracting a synthesised circle from images with different levels of additive noise. Ten different initialisations were used for each of the Kass technique, the dual gradient descent technique and the new search technique, and applied to images with six different levels of additive noise. One of these initialisations is shown in Figure 3(a) for which the noise standard deviation was 1.2, causing the target circle to be submerged within the noise. The other levels of noise were 0 (no noise), 0.4, 0.8, 1.2 and 2.0 (where the target circle was very difficult to perceive amidst the noise). The two contours in Figure 3(a) were used to initialise the search-based and gradient descent dual contours, the outer contour was used to initialise the Kass snake. Figure 3(b) illustrates the traditional problems with the Kass formulation. The contour can be driven over the features of interest, leading to a contour which does not match the target feature well. One single portion has reached the target boundary, the remainder is focused on noise. A further problem is parameterisation: it is possible to select value for parameters which would lead to a different result. A smaller value of the time step can allow the technique not to be driven over features of interest, but it may fail to reach the target feature and operation becomes slower. Alternative values for the

continuity, curvature, and image weighting parameters can alter the final result achieved, and with no statistics to guide their choice, together with the implicit scale dependent contraction force, successful application of the technique can be difficult. Alternatively, Figure 3(c) shows that the dual gradient descent contour can converge to the correct target. This is in part because the initialisation constrains the solution to lie between the inner and the outer contour. Also, difficulty in parameterisation is reduced since only a value for the regularisation parameter needs to be chosen, as opposed to values for the continuity, curvature, and image weighting parameters. Both gradient descent techniques require a termination parameter which is not required in the search-based technique. Finally, Figure 3(d) shows the result of the search-based technique. The regularisation parameter was chosen as  $\lambda = 0.5$  for the results in Figures 3(c) and 3(d). Though the final contour delivered by the search-based technique locates the position of the circle precisely, the final contour is actually a poorer match to the circle's perimeter than the result achieved by the dual gradient descent approach. This is due to the discretisation of the search space (the selection of the value for  $M$ ). Clearly,  $M$  could have been larger, giving finer resolution, but this would incur an increase in computation. As  $M$  is increased beyond the image resolution, there is a point where the global minimum is located correctly. Increasing  $M$  further will provide a more accurate localisation of the minimum, but for computational reasons it may be prudent to adopt a hybrid approach and switch to using a gradient descent technique, since the initialisation problem has now been solved. This hybrid approach would use the result of a search-based technique to initialise a local gradient descent technique. The performance advantage of

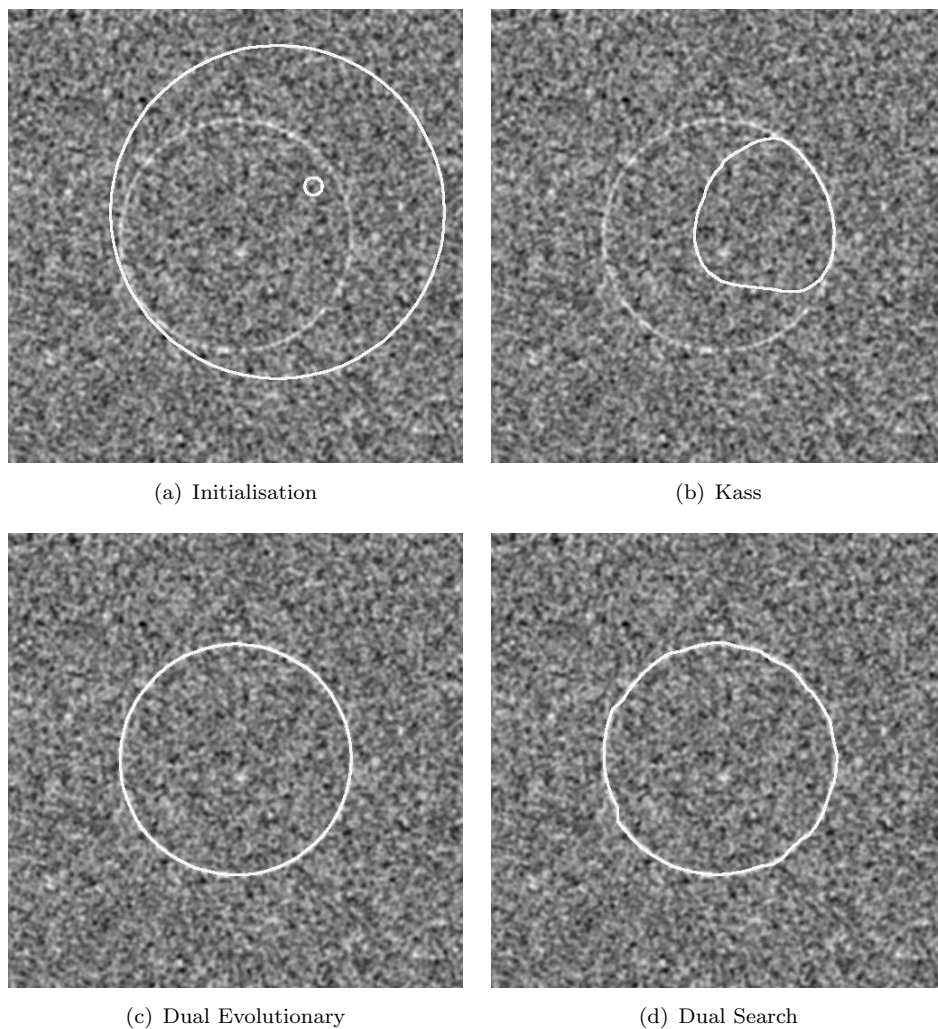


FIGURE 3: Example Results on a Synthetic Image

the search-based technique is revealed by analysis of the overall performance of the three techniques on synthetic images, over all noise levels and initialisations. To compare the final contour,  $C$ , with the target feature (the circle),  $T$ , we use the difference,  $d$  between their Fourier descriptors,

$$d(C, T) = \sqrt{\sum_{n=-F}^F |C_n - T_n|^2}, \quad (7)$$

where  $F$  is the number of descriptors. This measure is equal to the mean square deviation between the two contours in the spatial domain, but computation in the frequency domain is much more efficient. The measure varies with position, size and rotation, as required when matching to a circle with fixed position and size. The overall results are shown in Figure 4 as the average error plotted against increasing noise standard deviation. The average error is the difference  $d$  between the final contour and the target circle averaged over the ten different initialisations. The error in the final contour obtained by the Kass technique increases most rapidly with increase in noise standard deviation. The error in the results achieved by the dual gradient descent approach is better than that achieved by the Kass technique and becomes worse than the search-based dual contour when the noise standard deviation exceeds 1.2. The search-based technique performs best and its error only starts to increase markedly when the noise standard deviation exceeds 1.6. Note that the error in the search-based technique is slightly larger for lower levels of additive noise, consistent with its constraint to lie within the discrete search space (as shown in Figure 3(d)). The improved performance of the search-based contour at higher noise levels arises from its lack of prejudice to solutions near the initialisation as with gradient descent based techniques, but considers all solutions within the initialisation region. Furthermore it avoids the difficult problem of determining the evolution parameters, and requires a single regularisation parameter. These simulation results serve to illustrate the differing characteristics




FIGURE 4: Performance of Different Approaches with Noise

of local and global minimisation techniques. Analysis of application to real images is essentially application specific, requiring subjective appraisal of the results achieved as opposed to the objective appraisal of simulation results. Accordingly, we have used the new search-based strategy to determine the inner contour of a human face, an application where its properties are particularly required.

### 3 Head Boundary Extraction

Active contours are an attractive choice for the extraction of face boundaries. They appear well-suited, since face boundaries are not of uniform shape and exhibit low overall curvature. Location and description of these boundaries is important in facial feature extraction (Huang and Chen, 1992) and in model-based coding (Welsh et al., 1990). The inner face boundary (the border of the chin, cheeks and lower hair-line) has also been used in face recognition studies, as part of an extended feature vector (Jia and Nixon, 1995). For the face boundary, classical techniques based on local edge data are limited by their inability to provide an implicit description and can fail to exploit the inherent continuity of the face boundary, unlike active contours. Waite and Welsh (1990), Lam and H.Yan (1994) and Huang and Chen (1992) have all used snakes, using a closed contour, to extract the head boundary. However the extracted boundaries differ; Waite extracts a boundary including the chin and the upper part of the hair using an external contracting contour, whereas Huang and Lam extract the boundary of the chin and lower part of the hair using an internal expanding contour. Lam takes account of the actual face, by heuristically varying the snake parameters around the face boundary. Lam and Huang both use the greedy method (Williams and Shah, 1992) for the active contour technique. This technique is fast, but is restricted by the inability to guarantee a local minimum solution. All three approaches are limited by using a local minimisation approach and hence they require an unobstructed initialisation to enable the target feature to be extracted correctly. To illustrate the advantages that can be accrued

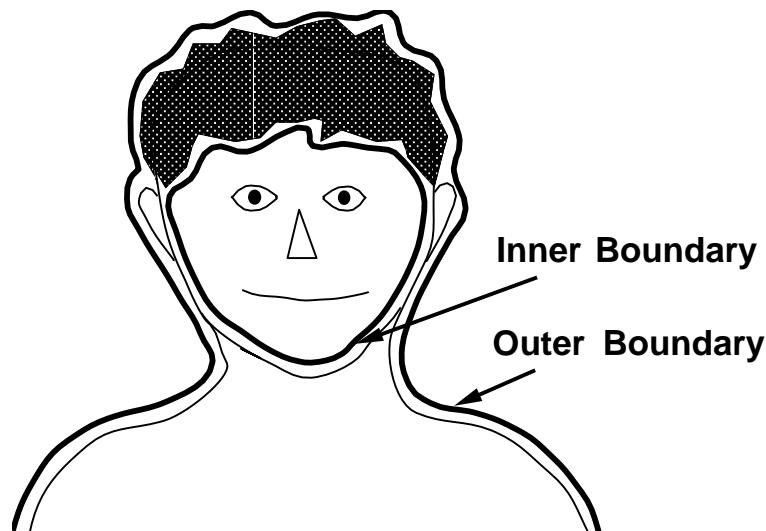


FIGURE 5: Head Boundary Definitions

by our new formulation, it is applied to extract the head and face boundary in front-view face images. Our approach uses two contour extraction techniques: an open contour to extract the outer head and shoulders' boundary and a closed contour to extract an inner boundary of the chin and lower hair line, Figure 5. In order to extract the inner and outer head boundaries two different active contour approaches are employed. The outer contour is found by a local gradient descent approach where the local nature is exploited by assuming the face to be on a constant background. The inner contour is found by the new global search-based technique, where its advantages are suited to a cluttered target (the inner face boundary) where an unobstructed initialisation is unrealistic. By application to a face database of 75 images, we show how this new technique can extract the boundaries of interest successfully.

Implementation of these techniques assumes a front-view face image on a plain background. An initial estimate of the head boundary is required to prime the gradient descent normal-driven snake used to

extract the outer contour (Cohen and Cohen, 1993). The inner contour extraction uses the results of the outer contour to calculate a suitable search space for the initialisation for the inner search-based technique. Alternative methods for locating the head boundary from a complex background include, a difference image exploiting temporal properties (Turk and Pentland, 1991), and neural networks (Sung and Poggio, 1995) searching to locate facial features. In this study, the emphasis is placed on the extraction technique. Accordingly a simple location scheme based on a binary edge image is employed such as would be available from a temporal difference-image approach, although in our experiments an edge functional is used since a plain background is assumed. A convex hull is used to extract a suitable contour representation to initialise the outer contour.

### 3.1 Outer Boundary Extraction

The extraction of the outer contour makes the assumption that the head is on a plain background. Consequently any edge data must lie on the boundary or within the head. The original head images are intensity normalised, filtered using a  $9 \times 9$  Gaussian mask ( $\sigma = 1.0$ ) and then passed through a first-order gradient-based edge operator. The resulting image is then thresholded to obtain a binary edge image. A fixed threshold value was used, which was selected to extract those parts of the edge image with intensity above the edge noise floor. The initial contour is obtained by generating an open convex hull. The end points are initially determined by searching with a vertically scanning technique, as in Figure 6. The resulting convex hull provides the initialisation for an open contour. This contour



FIGURE 6: Convex Hull Extraction

is then minimised according to a gradient descent based strategy. A normal force is applied to push the contour towards the head boundary and allow the contour to find non-convex solutions. The normal force parameter was determined heuristically on a limited selection of images, and was then fixed for the remainder of the tests.

### 3.2 Inner Boundary Extraction

The inner boundary is more difficult to extract with a local minimum based technique due to the quantity of edge data in the face boundary region. To develop a robust extraction technique it is necessary to use a search based strategy which overcomes the sensitivity to initialisation of the gradient descent techniques. Furthermore a particular problem in the extraction of head boundaries

is that the chin is often poorly characterised in the edge functional as a consequence of poor image contrast. Normal driven techniques will often drive the snake out of the weak minimum associated with the chin. The search-based technique has a greater tolerance with respect to initialisation than a comparable gradient descent technique, as reflected in the simulation results. Simple geometrical reasoning about the outer boundary is used to provide an approximate initialisation, Figure 7. The search space uses the centroid of the outer contour as its origin. The outer part of the search region is comprised of the top half of the outer contour augmented with a semi-ellipse to complete the bottom half. The semi-ellipse has a fixed aspect ratio of 1.5. The inner region is defined by a circular contour, of radius equal to half the minor axis of the ellipse. The circle is vertically offset by a constant factor. A constant value of  $\lambda = 0.7$  was used for the regularisation parameter, making no distinction between different parts of the face boundary. The number of snake points was fixed at  $N = 64$ , and the constraint lines were subdivided into  $M = 40$  points.




FIGURE 7: Inner Contour Search Space

## 4 Results

The head boundary extraction technique was applied to the face database taken during an undergraduate photo session and consequently the majority of subjects are trying to smile and have prepared hair styles. The lighting was controlled, and illuminated the subjects from the front. The database included male, female, bearded, spectacle wearers and two subjects wearing head gear. The initial contours for both outer and inner boundary extraction were assessed for correctness as follows: a correct outer initialisation is exterior to the head and terminates at the shoulders, and a correct inner initialisation describes an annular region containing the inner boundary. The results are given in Table 1. Failures in the outer initialisation were mainly due to poor contrast in the hair region. As a

Initialisation	Outer	Inner
Correct	71	55
Incorrect	4	20

TABLE 1: Initialisation Results

result there were no edges marked in the binary edge image causing the convex hull to cut across the head. The failures in the inner initialisation were partly due to the dependence on the outer initialisation and extraction result, but primarily on the weakness of the geometrical model used to derive the inner contours. Wide faces produced an over-sized interior contour which made the extraction of the correct contour impossible. The outer contour then produced a large region dropping below the neckline.

For assessment of contour extraction, the inner and outer contour results were overlaid on the original image and the correspondences were assessed. The results were rated subjectively according to accuracy as: "excellent", "good" or "poor". "Excellent" implies that the result could not be improved by manual intervention, whereas a "good" result could, but only for a very small region. A "poor" result implies that the result is obviously wrong in more than one respect. Alternatively, an empirical measure, such as described in (Gunn and Nixon, 1994) or one based on Equation 7, could be employed to interpret the results. The results for the contour extraction are shown in Table 2. The results for

Boundary	Outer	Inner
Excellent	37	38
Good	33	12
Poor	5	25

TABLE 2: Head Boundary Contour Results

six example results are shown in Figure 8, with (from top to bottom): outer initialisation, outer result, inner initialisation and inner result. Three of the "excellent" results are Faces 1, 2 and 3 which are shown in Figures Figure 8(a), Figure 8(b) and Figure 8(c). In these, the convex hull can be seen to lie on the top of the head and terminate correctly at the shoulders. This primed the open contour well and the outer contour can be seen to track the upper head and shoulders' boundary correctly. This gave a good initialisation for the dynamic programming technique, with an appropriate inner contour. The final inner contour extraction for both faces can be seen to track the inner face boundary correctly. Three faces showing reduced performance are also shown: Faces 4, 5 and 6 in Figures Figure 8(d), Figure 8(e) and Figure 8(f). The outer contour for Faces 4 and 5 were both rated only as "good" since the left shoulder is slightly incorrect for Face 4 and a small part of the left ear was missed in Face 5. Both contours are "excellent" in all other regions. The outer contour for Face 6 was rated as "poor" since part of the ear is missed, and the contour adheres to background shadow, on the right hand side of the face. The inner contours of Faces 4, 5 and 6 were rated as "poor": Face 4 because the shadow on the right side of the cheek is more prominent than neighbouring boundary causing the final contour to include this section erroneously; Face 5 because the contour incorrectly describes the hairline; Face 6 because the contour is attracted to strong edges in the eye region. These results highlight a weakness in the technique whereby the global minimum does not correspond to the full face boundary and could only be compensated for with additional prior information. However, the chin is extracted precisely in all the presented results, even though this part of the boundary it is often poorly defined in the edge data, reflecting the enhanced performance available from a global approach. The overall performance is summarised in Table 3. The extraction results were computed by considering the "excellent" and "good" as acceptable solutions. To determine the performance of the snake extraction techniques in isolation, the third column gives the results only for those in which the initialisation was "correct". The outer contour was extracted with a 96% success rate, and the inner contour with an 82% success rate. Inclusion of face features within the extraction technique appears eminently suitable as an approach to improve the initialisation of the inner search region.

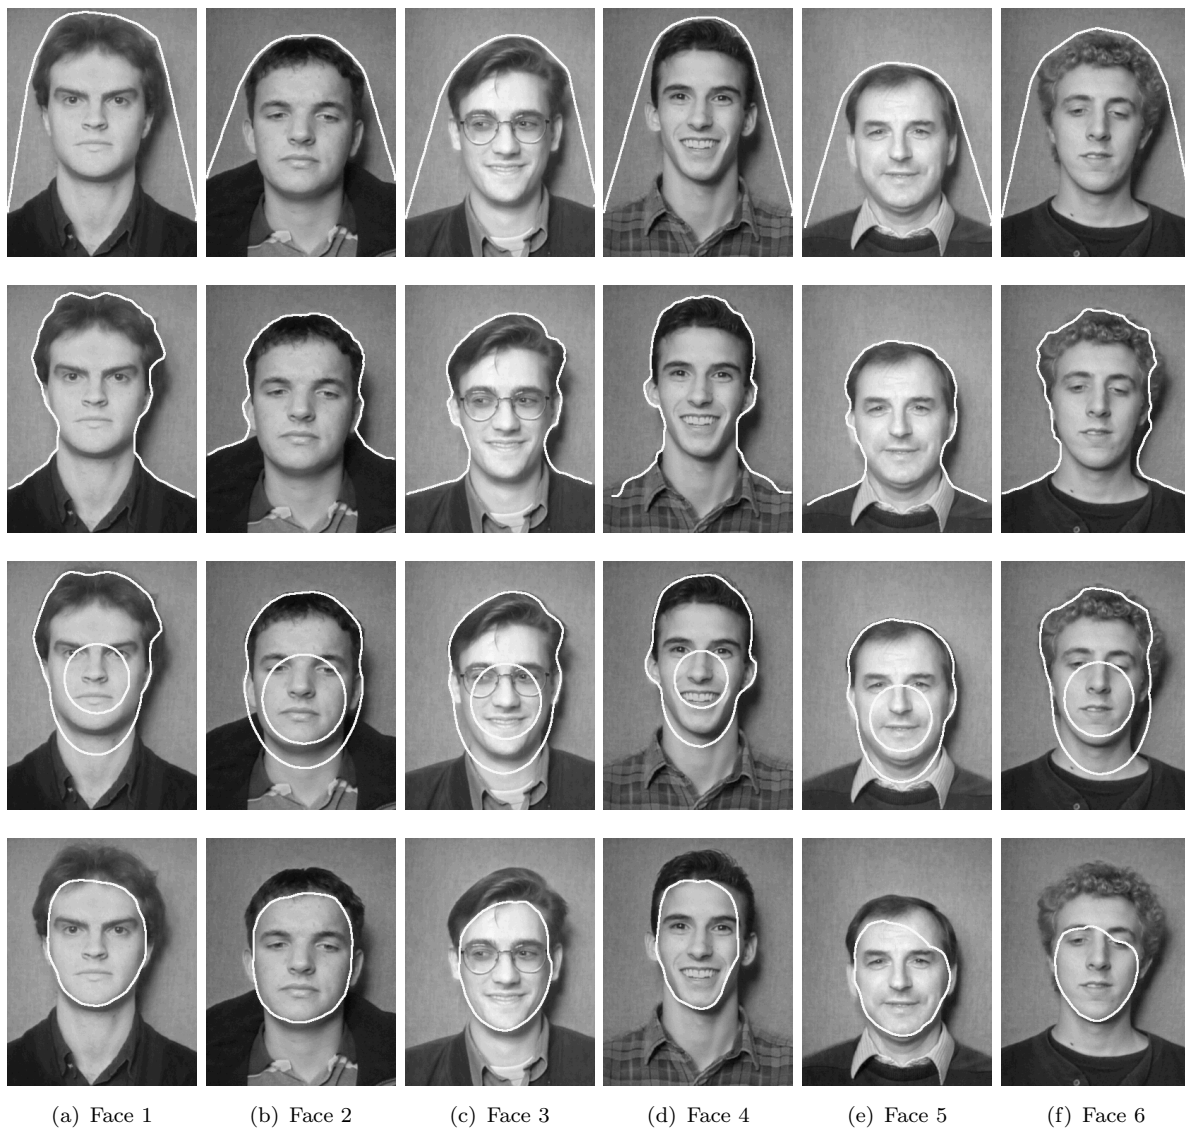


FIGURE 8: Example Head Boundary Extraction Results

	Initialisation	Extraction	Extraction given correct Initialisation
Outer	95%	93%	96%
Inner	73%	67%	82%

TABLE 3: Performance Results

## 5 Boundary Recognition

The extracted contour descriptions provide a method for assessing the similarity of face boundaries. To assess this possibility the internal boundaries from the previous experiment were used. They were chosen in preference to the outer boundary because the inner boundary contains less disturbance from time-variant features, such as the hair. The results that were classed as "excellent" from the inner contour extraction were assembled to form a known-face contour database. Then two 'unknown' inner contours were obtained by rerunning the contour extraction technique on two images with different pose, position and facial expression. These unknown boundaries were then compared against the 38 'known' contours to find a best match. The dissimilarity measure was computed as a normalised difference,  $nd$ , between two contours  $C$  and  $T$  (again expressed in Fourier descriptors) (Persoon and Fu, 1977):

$$d(C, T) = \min_{s, \varphi, \alpha} \sqrt{\sum_{n=-F}^{F, n \neq 0} |C_n - se^{j(n\alpha + \varphi)} T_n|^2}, \quad (8)$$

where the suffices denote the order of the descriptors,  $F = 20$  is the number of descriptors,  $s$  is the scale,  $\varphi$  is the orientation and  $\alpha$  is the starting point. Equation 8 measures the difference between two contours as in Equation 7, except that any differences in translation, scale, rotation and starting point are removed (the difference is invariant to translation by omission of the zeroth-order descriptor). The

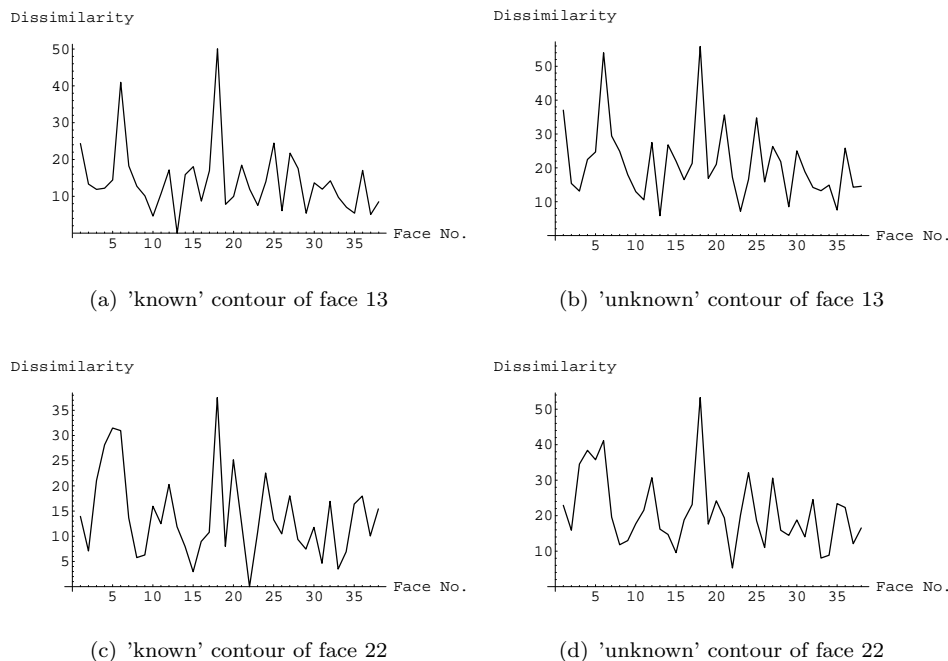


FIGURE 9: Similarity graphs

dissimilarity graphs in Figure 9 demonstrate the discriminatory ability of the inner face contour and show the dissimilarity  $nd$  when comparing the known and unknown face contours with the database of 38 known contours. Figures 9(a) and 9(c) illustrate how the 'known' contours of faces 13 and 22 compare with the other 'known' contours, and give an indication of the variance. In Figure 10(a), Face 13 is most similar to itself, and in Figure 9(c) Face 22 is most similar to itself, but the difference is greater for all other members of the database, as expected. With the chosen metric, Face 10 and Face 29 have a similarity to Face 13, and Faces 15 and 33 have a similarity to Face 22. Figures 9(b) and 9(d) show how the 'unknown' contours of faces 13 and 22 compare to the known database. It

can be seen that the technique correctly identifies both faces 13 and 22 as the best match. However, the measure is no longer zero in both cases, giving a dissimilarity of about 5. Figures 10(a), 10(c), 10(b) and 10(d) illustrate the extracted inner contours used for the unknown and known contours of Faces 13 and 22. The initial contour in Face 13 is slightly rotated from its known version, but a more significant difference concerns the upper right portion of the face where the contour tracks the hair-line better in the 'unknown' contour. The contour of Face 22 is larger in the unknown version, as compared with the known version. This is accommodated by the scale normalisation in Equation 8. Figures 10(e) and 10(f) show the two contours that were the most dissimilar to faces 13 and 22, though this difference is hard to perceive visually.




FIGURE 10: Inner extracted boundaries

Figure 11 gives a confusion matrix that expresses the match of all known contours to themselves; a black dot represents a good match, whereas a white dot represents a poor match. As such, Figures 9(a) and 9(c) are the horizontal slices through the confusion matrix for faces 13 and 22, respectively, where the minima in Figures 9(a) and 9(c) correspond to the dark points in Figure 11. Figure 11 clearly shows the difference between the known contours and the face database.




FIGURE 11: Confusion Matrix by Matching Face Contours

## 6 Conclusions

Implementations of active contour techniques can use gradient descent or global search to determine energy minima in feature extraction. The concept of a dual contour approach has been retained but minimisation has been replaced with dynamic programming to search for a global contour solution within the target region. Extraction by dynamic programming is superior because only a single regularisation parameter is required. Gradient descent techniques require additional parameters such as a termination parameter controlling when evolution is halted. The two techniques have been shown to provide complementary characteristics for contour extraction. Simulation results have demonstrated that the search-based approach can provide a better solution in noisy imagery, consistent with lack of availability of an unobstructed initialisation. Since head boundaries are of low overall curvature, and an unobstructed initialisation is unrealistic for the inner face boundary, the new search-based approach gives an appropriate basis for head and face boundary extraction. The new technique has been applied to a database of 75 face images and resulted in a successful extraction rate of 96% for the outer contour and 82% for the inner contour, given a "correct" initialisation. The results showed the need to improve the initialisation techniques for a more robust head boundary technique for, say, a face recognition system. The results demonstrates that with valid initialisations both the gradient descent and search-based active contour technique provide a promising approach for outlining and boundary extraction respectively. The main limitation is the dependence upon the quality of the initialisation, and the quality of the edge detection, which is affected by illumination and boundary contrast.

## References

- A.A. Amini, T.E. Weymouth, and R.C. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(9):855–867, 1990.
- L.D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2d and 3d images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1131–1147, 1993.
- D. Geiger, A. Gupta, L.A. Costa, and J. Vlontzos. Dynamical programming for detecting, tracking and matching deformable contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(3):294–302, 1995.
- S.R. Gunn and M.S. Nixon. **A model based dual active contour**. In E. Hancock, editor, *Proc. British Machine Vision Conference*, pages 305–314, York, U.K., 1994. BMVA Press.
- S.R. Gunn and M.S. Nixon. **A robust snake implementation: A dual active contour**. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19(1):63–68, 1997.
- C.L. Huang and C.W. Chen. Human facial feature extraction for face interpretation and recognition. *Pattern Recognition*, 25(12):1435–1444, 1992.
- X. Jia and M.S. Nixon. Extending the feature vector for automatic face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(12):1167–1176, 1995.
- M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1988.
- K.F. Lai and R.T. Chin. Deformable contours - modeling and extraction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(11):1084–1090, 1995.
- K.M. Lam and H. Yan. Locating head boundary by snakes. In *Proc. Int. Symposium on Speech, Image Processing and Neural Networks*, Hong Kong, April 1994.
- S. Menet, P. Saint-Marc, and G. Medioni. Active contour models: Overview, implementation and applications. In *Proc. IEEE International Conference on Systems, Man, And Cybernetics*, volume 212, pages 194–199, 1990.
- E. Persoon and K.S. Fu. Shape discrimination using fourier descriptors. *IEEE Transactions on Systems, Man, and Cybernetics*, 7(3):170–179, 1977.
- D. Rueckert and P. Burger. Contour fitting using an adaptive spline model. In *Proc. British Machine Vision Conference*, pages 207–216, Birmingham, U.K., 1995.
- G. Storvik. A bayesian approach to dynamic contours through stochastic sampling and simulated annealing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):976–986, 1994.
- K.K. Sung and T. Poggio. Learning human face detection in cluttered scenes. In *Proc. Computer Analysis of Images and Patterns*, pages 432–439, Prague, 1995.
- M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.
- J.B. Waite and W.J. Welsh. Head boundary location using snakes. *British Telecom Technology Journal*, 8(3):127–136, 1990.

- 
- W.J. Welsh, S. Searby, and J.B. Waite. Model-based image coding. *British Telecom Technology Journal*, 8(3):94–106, 1990.
- D.J. Williams and M. Shah. A fast algorithm for active contours and curvature estimation. *CVGIP: Image Understanding*, 55(1):14–26, 1992.
- G. Xu, E. Segawa, and S. Tsuji. Robust active contours with insensitive parameters. *Pattern Recognition*, 27(7):879–884, 1994.