Computer Science Journal of Moldova, vol.**14**, no.3(42), 2006

Applying Evolution Strategies for Chest Radiographs Segmentation

Oliviu Matei

Abstract

Segmentation is one of the key areas in computer vision. Advanced techniques, such as active contour models and active shape models suffer from many drawbacks. An improvement of the two approaches, based in evolution strategies, is presented. Some of its advantages lie in less computational resources and less a priori knowledge required. At the end of the paper, the results of the experiments carried out with the new approach are presented.

1 Introduction

The ultimate goal of machine vision is image understanding - the ability not only to recover image structure, but also know what it represents. An attractive field of application for image segmentation is the medicine. Almost every object of interest in the human body can vary in size, shape and appearance. Often this makes the task of automatically identifying and segmenting interesting structures, such as organs and bones, very difficult.

This article presents a segmentation technique based on evolution strategies. We compare this method with the classical ones - active contour models and active shape models. At the end we show the results of image segmentation applying the proposed technique.

1.1 Related work

Many researchers approached the segmentation problem in a bottom-up fashion: emphasis was on the analysis and design of filters for the detection of local structures such as edges, ridges, corners and 'T'-junctions. The structure of an image can be described as a collection of such syntactical elements and their (spatial) relations, and such descriptions can be used as input for generic segmentation schemes. Unfortunately, these segmentations are often not very meaningful. On the other hand, top-down strategies (also referred to as model-based or active approaches) for segmentation were used successfully in highly constrained environments, e.g., in industrial inspection tasks. Often these methods are based on template matching. But template matching, or related techniques, are likely to fail if the object and/or background exhibit a large variability in shape or gray-level appearance, as is often the case in real-life images and medical data. Active contours or snakes [3], [9] and wave propagation methods such as level sets [16], have been heralded as a new paradigms for segmentation.

Other researchers experimented with hand-crafted parametric models. An illustrative example is the work of Yuille et al. [17] where a deformable model of an eye is constructed from circles and parabolic patches and a heuristic cost function is proposed for the gray-level appearance of the image inside and on the border of these patches.

There is a need for generic segmentation schemes that can be trained with examples as to acquire a model of the shape of the object to be segmented (with its variability) and the gray-level appearance of the object in the image (with its variability). Such methods use statistical techniques to extract the major variations from the prototypes in a principled manner. Several of such schemes have been proposed. One of the methods is active shape models (ASMs) put forward by Cootes and Taylor [7]. The shape model in ASMs is given by the principal components of vectors of landmark points.

2 Statistical Shape Models

Here we describe a statistical model of shape that will be used to represent objects in images. The shape of an object is represented by a set of points. Our aim is to derive models which allow us to both analyze new shapes and to synthesize shapes similar to those in the training set.

Good choices for landmarks are points that can be consistently located from one image to another. In two dimensions, points could be placed at clear corners of object boundaries, 'T'-junctions between boundaries or easily located biological landmarks. However, there are rarely enough of such points to give more than a sparse description of the shape of the target object. This list would be augmented with points along boundaries that are arranged to be equally spaced between well defined landmark points.

3 Active Contour Models

An active contour (also known as a snake by virtue of the nature of its evolution) is an energy minimizing spline ([4]) that is constrained by its own internal forces of continuity and curvature, while external forces drive it towards desired image features. The external forces are provided by the image and where appropriate, information from a higher-level process may be included. The problem is formulated as an energy minimization problem. Typically there are many local minima in the energy function. In this situation the contour may become trapped in strong neighboring edges resulting in a false boundary. Prior knowledge may be included in the process in order to achieve this. A model-free interpretation is limited because the problem is underconstrained. Various methods have been proposed to avoid insignificant false edges in [3], [8] and [14]. However, even with the application of these techniques the contour may still become trapped by false edges.

There are two key difficulties with active contour algorithms. First, the initial contour must, in general, be close to the true boundary or else it will likely converge to the wrong result. The second problem is

that active contours have difficulties progressing into concave boundary regions [17].

4 Active Shape Models

This section briefly reviews the ASM segmentation scheme. We follow the description and notation of [6]. In principle, the scheme can be used in nD, but in this paper we give a two-dimensional (2-D) formulation.

A shape model relies on representing the objects by a set of labelled points; each point is placed on a particular part of the object. The model gives the average positions of the points and a description of the main modes of variations found in the training set. Locating an example of such a model in an image involves choosing values for each of the parameters so as to best fit the model to the image. An initial guess for the best shape, orientation, scale and position can be refined by comparing the hypothesized model example with the image data and using differences between model and image to deform the shape. The method has similarities with the ACM's, but differs in that global shape constraints are applied; to make this distinction clear, the term "Active Shape Models" was adopted. The key point is that an instance of a model can only deform in ways found in its training set.

5 Evolution Strategies

Evolution strategies (ESs) are search algorithms which imitate the principles of natural evolution as a method to solve parameter optimization problems. They are based on the idea that among a population of examples, only members which are the fittest can survive and breed so that the examples in following generations are improved.

A general evolutionary algorithm is defined below:

Generate random population
REPEAT

evaluate fitness of current population
select chromosomes, based on fitness
for reproduction perform crossover and mutation
to give new improved population

UNTIL finished

An evolution strategy must have the following five components:

1. a definition of the objective function
2. a definition and implementation of the genetic representation
3. a way to create an initial population of potential solutions
4. a definition and implementation of the genetic operators
5. values for various parameters that the evolution strategy uses (population size, probabilities of applying operators etc.)

Once these three have been defined, the evolution strategy should work fairly well. Beyond that, one can try many different variations to improve performance, find multiple optima (species - if they exist), or parallelize the algorithms.

5.1 Objective Function

In our experiments, we applied three types of objective functions:

- The first function used was the magnitude of the edges. We assumed that the object (specifically the lung) is delimited by clear edges.

  The magnitude (or edge strength) is computed using the Sobel operator. It uses a pair of 3x3 convolution masks, one estimating the gradient in the x-direction (columns) and the other estimating the gradient in the y-direction (rows).

$$G_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} G_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

The magnitude is then:

$$|G| = |G_x| + |G_y|$$

Since the edges are defined as the pixels with the highest magnitude, the evolutionary strategy has to find those pixels which maximize the magnitude. Mathematically, this is:

$$f(x) = \sum_{i=1}^{n} G_{p(i)}$$

where $f(x)$ is the objective function and $G_{p(i)}$ is the magnitude of the pixel $p(i)$.

This was performed in two ways.

The first approach was the detection of edges in a raw image, respectively in a smoothed image. As expected, the results were much better in the smoothed image. On the raw image, the population got stuck often in local optima. Whereas a smoothing of the image assured a "smoother" intensity of the pixels. Another approach was based on edge detection on a sharped image. The sharpening was performed using Laplace operator. A disadvantage of this operator is the doubling of some edges in the image and the small neighborhood it takes into account. To prevent these, we applied a gaussian operator.

Compared with the previous approach, the results became better. Sharpening the image, the search was not performed on the image space (formed by the intensities of the pixels), but rather on the gradient space, which is the absolute difference between neighboring pixels. This makes the approach less sensitive to intensity variations.

- The second objective function was an extension of the previous one. For sharpening the image, we took into account not only the magnitude, but also the direction of the edge. The sharpening was performed using Prewitt operator.

In this case, for each landmark there were two criteria: the magnitude and the direction of the edge in that point. The direction of an edge is determined using the Susan edge detector čitesmith99susan, which, because of the lack of space, will not be detailed here.

In this case, the objective function is:

$$f(x) = \alpha \sum_{i=1}^{n} G_{p(i)} + \beta \sum_{i=1}^{n} R_{p(i)}$$

where $G_{p(i)}$ is the magnitude of the pixel $p(i)$, $R_{p(i)}$ is the direction of the edge in the pixel $p(i)$ and α and β are two parameters tuning the importance of one factor or the other. They must fulfill the condition: $\alpha + \beta = 1$. If $\alpha = 1$, then this objective function is identical with the previous one. If $\alpha = 0$, then the objective function takes into account only the direction of the edges.

This approach runs well only on some images. That is because corresponding landmarks do not have corresponding edges directions. However, on shapes that can be easily defined mathematically, the results were very satisfactory.

- The third function is a more general case, in which we consider not only the magnitude of the edge, but also the gray-level appearance of the edge.

  In some cases it is efficient to assume that the points lie on strong edges and to search for such in an image. However this is not always satisfactory and it is necessary to have more general model of the gray-level appearance. The gray-level appearance model that describes the typical image structure around each landmark is obtained from pixel profiles, sampled (using linear interpolation) around each landmark, perpendicular to the contour.

  This requires a notion of connectivity between the landmark points from which the perpendicular direction can be computed. The direction perpendicular to a landmark (x_n, y_n) is computed

by rotating the vector that runs from (x_{n-1}, y_{n-1}) to (x_{n+1}, y_{n+1}) over 90°. We assume that all objects we use are closed contours, so for the first landmark, the last landmark and the second landmark are the points from which a perpendicular direction is computed; for the last landmark, the second to last landmark and the first landmark are used.

On either side k pixels are sampled using a fixed step size, which gives profiles of length $2k + 1$. Cootes and Taylor [6] propose to use the normalized first derivatives of these profiles to build the gray-level appearance model. The derivatives are computed using finite differences between the $(j-1)$th and the $(j+1)$th point:

$$g_{ijk} = I_j(y_{i(k+1)}) - I_j(y_{i(k-1)}) \tag{1}$$

where y_{ik} is the kth point along the ith profile and $I_j(y_{ik})$ is the gray level in image j at that point. The normalization is such that the sum of absolute values of the elements in the derivative profile is 1:

$$g'_{ij} = \frac{g_i j}{\sum_{k=1}^{n_p} |g_{ijk}|} \tag{2}$$

Denoting these normalized derivative profiles as $g_1, g_2, ..., g_s$, the mean profile $\overline{g}$ and the covariance matrix S_g are computed for each landmark.

$$\overline{g_i} = \frac{1}{N_s} \sum_{j=1}^{Ns} g'_{ij} \tag{3}$$

This allows for the computation of the Mahalanobis distance [11] between a new profile and the profile model

$$f(g_i) = (g_i - \overline{g}) S_g^{-1} (g_i - \overline{g}) \tag{4}$$

Minimizing $f(g_i)$ is equivalent to maximizing the probability that g_i originates from a multidimensional Gaussian distribution. In other words, minimizing $f(g_i)$ is equivalent to maximizing the probability that the points are the desired ones.

5.2 Genetic Representation

An individual is represented by all the landmark points for a training image. Noting each landmark point by (x_i, y_i), the set of all landmarks is then:

$$x = (x_1, y_1, x_2, y_2, ..., x_n, y_n) \tag{5}$$

where n is the number of landmarks defined for one training image.

However, an interesting idea introduced in ES is that an individual is not a simple float or binary value, like in the case of GAs, but is a pair of float values, i.e., $v = (x, \sigma)$, where the first vector x represents a point in the search space, and the second vector σ is a vector of standard deviations. For our case, an individual is represented by:

$$x = ((x_1, y_1, x_2, y_2, ..., x_n, y_n), (\sigma_1, \sigma_2, ..., \sigma_s, ..., \sigma_{2s})) \tag{6}$$

5.3 The Initial Population

Once a suitable representation has been decided upon for the chromosomes, it is necessary to create an initial population to serve as the starting point for the genetic algorithm. The initial population is usually created randomly. When there is some information about the solution, some heuristics can be used. Each subsequent population that is built (each subsequent generation) uses the previous population as a base - taking the more fit individuals to breed better solutions.

In this case, we can use the training landmarks as forming the initial population, rather than random points. However, a random population should lead to the same good results in segmentation, although after a larger number of generation. On the other hand, an heuristic initial population would yield good results faster because the initial shapes resemble the one to be segmented.

5.4 Genetic Operators

To produce an offspring, the system acts in several stages:

1. Select two individuals:

$$(x^1, \sigma^1) = [(x_1^1, ..., x_n^1, y_1^1, ..., y_n^1), (\sigma_1^1, ..., \sigma_{2n}^1)] \tag{7}$$

and

$$(x^2, \sigma^2) = [(x_1^2, ..., x_n^2, y_1^2, ..., y_n^2), (\sigma_1^2, ..., \sigma_{2n}^2)] \tag{8}$$

and apply a recombination operator. The types of crossover that we used are:

(a) discrete, in which case the offspring is

$$(x, \sigma) = [(x_1^{q_1}, ..., x_n^{q_s}, y_1^{q_{s+1}}, ..., x_{2n}^{q_{2n}}), (\sigma_1^{q_1}, ..., \sigma_{2n}^{q_{2n}})] \tag{9}$$

where $q_i = 1$ or $q_i = 2$ with equal probabilities for all $i = 1, 2, ..., s$. That means that each component i comes either from the first or from the second parent.

(b) intermediate, when the offspring is

$$(x, \sigma) = [(x_1^{q_1}, ..., x_n^{q_s}, y_1^{q_{s+1}}, ..., y_n^{q_{2n}}), (\sigma_1^{q_1}, ..., \sigma_{2n}^{q_{2n}})] \tag{10}$$

$$(x, \sigma) = [(\frac{x_1^1 + x_1^2}{2}, ..., \frac{x_n^1 + x_n^2}{2}, \frac{y_1^1 + y_1^2}{2}, ..., \frac{y_n^1 + y_n^2}{2}, (\frac{\sigma_1^1 + \sigma_1^2}{2}, ..., \frac{\sigma_{2n}^1 + \sigma_{2n}^2}{2})] \tag{11}$$

The intermediate crossover yielded better results. The explanation is somehow obvious. Whereas the discreet crossover maintains at least one of the coordinates of either parents, the intermediate crossover creates a completely new offspring, with completely new coordinates.

2. Apply mutation to the offspring obtained. Here is the novel feature of $(\mu + \lambda) - ES$ and $(\mu, \lambda) - ES$. By this operator, the

control parameters adapt themselves. After mutation, the individual $v = (x, \sigma)$ becomes $v' = (x', \sigma')$, where

$$\sigma' = \sigma \dot{e}^{N(0,\delta\sigma)} \tag{12}$$

and

$$x' = x + N(0, \sigma') \tag{13}$$

In the above relations, $\delta\sigma$ is an extra parameter and $N(0, \delta\sigma)$ represents a random number within a normal (Gaussian) distribution with mean zero and standard deviation $\delta\sigma$.

In ESs, mutation realizes a kind of hill-climbing search when it is considered in combination with selection. The σ_i parameters related to the x_i can be seen as preferred search directions along the axes of coordinates. Generally, the best search direction (with the largest climbing, in mathematics called *gradient*) is not along the axis. Thus the trajectory of the population through the search space is zigzagging along the gradient, hence it is suboptimal. In order to avoid this shortcoming, an additional control parameter (θ) has been introduced. An individual is represented by a triple (x, σ, θ). The recombination operators become now:

$$\sigma' = \sigma \dot{e}^{N(0,\delta\sigma)} \tag{14}$$

$$\theta' = \theta + N(0, \sigma') \tag{15}$$

$$x' = x + C(0, \sigma', \theta') \tag{16}$$

The parameter $\delta\theta$ is an extra parameter, like $\delta\sigma$. The parameter $C(0, \sigma', \theta')$ is the correlation matrix.

5.5 ES Parameters

An important step in designing an evolution strategy is the choice of the parameters.

1. The individual length is twice as much as the number of landmarks. If there are n landmarks (x_i, i_i) defined for each image, an individual is:

$$x = (x, \sigma, \theta) = [(x_1, ..., x_n, y_1, ..., y_n)(\sigma_1, ..., \sigma_{2n})(\theta_1, ..., \theta_{2n})] \tag{17}$$

2. The size of the population, μ represents the number of individuals that undergo evolution (crossover and mutation).

3. The number of offspring, λ is the number of individuals created after a generation evolves. It is an intermediary population, from which, the best individuals will survive and make up a new generation.

4. The selection procedure can be done in two ways:

   - selecting the fittest μ individuals from the intermediary population consisting of λ offspring. This type of selection is referred to as (μ, λ) selection
   - selecting the fittest μ individuals from the population consisting of the μ parents along with the λ offspring. This is referred to as $(\mu + \lambda)$ selection

5. Mutation parameters $\delta\sigma$ and $\delta\theta$ are needed for defining the step of the mutation. The larger the two parameters, the more significant the mutation.

6. The total number of generations represents the maximum number of iterations until the algorithm stops. Normally, the algorithm should end when the objective function is reached or when it is close enough. If this does not happens, the algorithm runs until the maximum number of generations.

6 Why Evolution Strategies?

6.1 Evolution Strategies vs. Genetic Algorithms

Both ESs and Genetic Algorithms (GAs) belong to a wider class of search algorithms, known as Evolutionary Algorithms or Evolutionary Computation. Both approaches to computational problem solving rely on the biological paradigms of organic evolution, postulated by Darwin. They explore the search space by means of population of search points which undergo evolution by crossover, mutation and selection.

Cootes et al. [5] and Huang et. al. [12] proposed the use of GAs for segmentation. Unlike them, we propose ESs for the same task. The reasons for that are described further.

1. A reason for choosing ESs consists in the way the individuals are represented. The ESs individuals are represented by floats. On the other hand, the GAs operate on binary strings. Only recently, they have been used also float representation.

2. The selection process represents another reason for using ESs. The new generation is selected deterministic as consisting of the fittest μ individuals out of λ (in the case of $(\mu, \lambda) - ES$) or $\mu + \lambda$ individuals (for $(\mu + \lambda) - ES$).

   In a generation of the GAs, all the individuals have a chance to be selected according to their fitness. Thus even the weakest ones can be parts of the new population. Thus the selection is weakly controlled.

3. Another reason for choosing the ESs lies in the way the reproduction parameters are stored and handled. In GAs, the probability of crossover and the probability of mutation remain constant throughout the evolution process and are external to the population. On the contrary, these parameters evolve every new generation in ESs. This self-adaptation is quite important for fine local tuning and has its counterpart in GAs research.

6.2 ESs vs. ACMs

As can be seen in section 3 the main disadvantage of the ACMs is their additional computational complexity. There are very complex calculations to be done, including integrals and derivatives.

Second, most approaches are based on explicit updating schemes which demand very small time steps. The steps used by ESs have two advantages:

- can be tuned with the evolution parameters, namely $\delta\sigma$ and $\delta\theta$

- vary over time, from small steps to larger ones, depending on the fitness of the individuals

Third, an ACM requires the definitions of both an external energy as well as an internal one. The external energy is dependent on the application and is similar to the target shape for ASMs and to the objective function for ESs. Straightforward functions have been defined by Gunn et. al. [10] and by Xu at al. [15].

However, the internal function may have significant drawbacks:

- It may be easily defined by regular shapes, but difficult or even impossible to define for more complex, real life shapes.
- Quite often only one function does not suffice. Most of the times, the internal forces of a shape can be defined by more functions.

ESs do not necessarily make use of the internal energies. However, such heuristic knowledge can be embedded successfully into the evolution, which leads to faster and more reliable convergence. In this case, the size of the population is important, because the individuals "pull" each other towards the contour of the object.

According to Yuille et al. [13], the active contours have difficulties progressing into concave boundary regions. On the other hand, the ESs do not have such difficulties as they do not make use of any function defining an internal energy.

6.3 ESs vs. ASM

Cootes et. al. [7] propose, as improvement of the ASMs, the alignment of images. This leads to better results. The drawbacks of such an approach are:

- It requires extra computational resources.
- The alignment works well on the training images, but the image to be segmented cannot be aligned because there are no landmarks defined yet. This makes that approach almost useless in real life.

The ASMs require an approximate position of the object to be segmented, whereas ESs are global optimization techniques. In other words, they "search" for the objects independently of the position and shape of the object on the training image. However, training shapes spread all over the images would yield better results in terms of time and segmentation.

The ASMs can get stuck in secondary edges and cannot be avoided due to the small steps in which the shape is updated. This situation can occur also in ESs. However, due to mutation, some individuals may get out of them and eventually reach the global optimum.

7 Results

We carried out three types of experiments. The objects in the images were annotated by a number of fixed landmarks and a closed contour between those fixed points from which a number of equidistant landmark points were sampled. No shape alignment was performed.

As shown in section 5.5, an individual consists of a set of landmarks describing a contour. A training set of images were marked manually by two physicians. An average of the two segmentations yielded the initial population, which, further on, underwent evolution.

To evaluate the results, the following measure was used: the area correctly classified as object was divided by the area classified as object plus the area classified incorrectly as object or as background. In other words: true positive area divided by true positive plus false positive plus false negative area. This performance measure yields 1 for a perfect result and 0 if there is no overlap at all between the detected and true object. This measure is probably more directly related to what is expected of a segmentation in practice than the average distance between the true and detected location of each landmark in the shape model because the latter is not sensitive to shifts of the landmarks along the contour.

The first set of experiments consisted of segmentation of rectangles of different sizes. This step was required for rough assessment of the

methods. It gave the first impressions of various setups. The number of fixed points was 4 (in the corners of the rectangles) and in total 24 points were used.

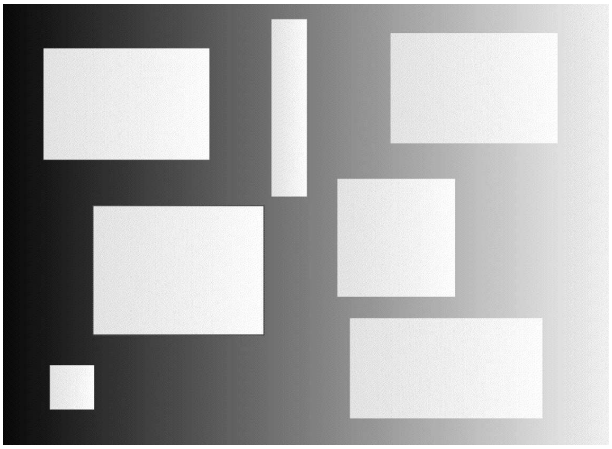

Figure 1. Rectangles used for segmentation

The other experiments were performed on 78 standard postero-anterior (PA) chest radiographs. In these images, both the left and the right lung fields were segmented as separate objects. The images were randomly selected from a tuberculosis screening program and contained both normal and abnormal cases of patients of 16 years and older. The images were printed on 10 x 10 cm film, digitized with a scanner to 1000 x 1000 pixels with 12 bit intensity.

The lung fields in the test set of the chest radiographs were manually segmented by a human observer independently. Therefore the results of the algorithms could be compared with the output of the observer. For these experiments, 15 fixed landmarks have been annotated and in total 60 points were used.

The second set of experiments consisted in segmentation of the left lung field.

The results we obtained are summarized in the Table 1.

Figure 2. Rx images used for left lung segmentation

Table 1. Segmentation results for the left lung fields

Segmentation	**Accuracy** $\pm\sigma$
Segmentation using the edge magnitude	$88.4\% \pm 7.4\%$
Segmentation using the edge magnitude and direction	$94.5\% \pm 9.87\%$
Segmentation using the gray-level appearance	$95.2\% \pm 6.4\%$

The third set of experiments consisted in segmentation of the right lung field.

Figure 3. Rx images used for right lung segmentation

The results we obtained for the right lung field segmentation are summarized in the Table 2.

As can be noticed from the experiments, the results for the left and for the right lung are pretty similar. The three segmentation criteria used yielded quite some different results. The edge magnitude gave the worst results, but with the smallest standard deviation. The second

Table 2. Segmentation results for the right lung fields

Segmentation	**Accuracy** $\pm\sigma$
Segmentation using the edge magnitude	$87.3\% \pm 5.21\%$
Segmentation using the edge magnitude and direction	$93.5\% \pm 8.6\%$
Segmentation using the gray-level appearance	$96.1\% \pm 6.58\%$

criterion, edge magnitude and direction resulted pretty good results, but the standard deviation was large, which means that it works well only on some images. The last criterion, gray-level appearance, gave the best results, with small standard deviation, which means good and constant results.

Further we will compare the results of our approach with results obtained by other authors. Bram van Ginneken et. al. [1] carried out similar research on lung fields segmentation. They used 230 standard chest radiographs and applied ASMs and ASMs with optimal features. Their results are concluded in the tables below. The results of the right lung segmentations are presented in Table 3.

Table 3. Segmentation results of the experiments carried out by Ginneken for the right lung fields

Segmentation	**Accuracy** $\pm\sigma$
ASMs	$88.2\% \pm 7.4\%$
ASMs with optimal features	$92.9\% \pm 2.6\%$

The results of the left lung segmentations are presented in Table 4.

It is difficult to compare results of two researches carried out on different sets of images and in different conditions. However, it is noticeable that the results presented in [1] and obtained by us are very similar. Moreover, unlike Ginneken, our results are more compact (σ is slightly lower in our case), which means more consistent segmentations.

We consider interesting to see the similarities and differences

Table 4. Segmentation results of the experiments carried out by Ginneken for the left lung fields

Segmentation	**Accuracy** $\pm\sigma$
ASMs	$86.1\% \pm 10.9\%$
ASMs with optimal features	$88.7\% \pm 11.4\%$

between segmentations based on evolutionary methods for various anatomical shapes. Huang et. al. [12] used genetic algorithms for eye location. They used 10 faces for training and 20 face images for tests, out of which their algorithm missed two eye locations and no false positives were identified. In other words, the accuracy of the algorithm was about 90%. Eye detection is a quite simple task because it involves a rather coarse segmentation.

Cootes et al [5] applied ASMs and genetic algorithms for locating structures in medical images. They defined a 96 point heart model on a set of echocardiograms. They obtained good results after 50 iterations, but did not assessed the accuracy of the location.

8 Conclusions

The method we describe allow flexible models of image objects such as lungs to be built easily from sets of example images. The technique can be applied to a wide variety of objects in different imaging modalities.

Cootes [5] proposed the use of genetic algorithms for the searches performed by the ASMs. Our approach, on the contrary, substitutes entirely the ASMs. The method we presented could be applied for segmentation of any anatomical structure. However, our experiments were carried out only for lung fields.

An enhancement of this approach is the use of the ESs in a multiresolution framework. This means the use of ESs within various resolutions, starting with a coarse image and ending with the finest resolution.

Beichel et al. [2] researched active appearance model matching.

A combination between active appearance with evolution strategies would be a very good research subject as well as an improvement of the approach we have presented in this paper.

References

[1] J. Staal B. ter Haar Romeny B. van Ginneken, A.F. Frangi and M.A. Viergever. Active shape model segmentation with optimal features. *IEEE Transactions on medical imaging*, 21(8), 2002.

[2] R. Beichel, H. Bischof, and M. Sonka. Robust active appearance model matching. In *Information Processing in Medical Imaging, 19th International Conference*, pages 114–125, 2005.

[3] D. L. Pham C. Xu and J. L. Prince. *Handbook of Medical Imaging, Vol. 2:Medical Image Processing and Analysis*. SPIE PRESS, 2000.

[4] Jr. C. Xu, A. Yezzi and J. L. Prince. On the relationship between parametric and geometric active contours. In *Proc. of 34th Asilomar Conference on Signals, Systems, and Computers*, pages 483 –489, 2000.

[5] T. F. Cootes, A. Hill, and J. Haslam. The use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–366, 2001.

[6] T. F. Cootes and C. J. Taylor. Statistical models of appearance for computer vision. Technical report, University of Manchester, 2001.

[7] T. F. Cootes, C. J. Taylor, D. Cooper, and J. Graham. Active shape models-their training and application. *Computer Vision Image Understanding*, 61(1):38–59, 2001.

[8] H. Delingette and J. Montagnat. Shape and topology constraints on parametric active contours. *Computer Vision and Image Understanding*, (83):140–171, 2001.

[9] D. Terzopoulos G. Hamarneh, T. McInerney. Deformable organisms for automatic medical image analysis. In *Proc. Third International Conference on Medical Image Computing and Computer Assisted Interventions*, volume 1, pages 66–75, 2003.

[10] Steve R. Gunn and Mark S. Nixon. A model based dual active contour. Technical report, University of Southampton, 2002.

[11] W. Haerdle, Y. Mori, and J. Symanzik. *Computational Statistics.* Springer Verlag, 2004.

[12] J. Huang and H. Wechsler. Eye location using genetic algorithm. In *2nd International Conference on Audio and Video-Based Biometric Person Authentication (AVBPA)*, 1999.

[13] C. English J. M. Coughlan, D. Snow and A.L. Yuille. Feature extraction from faces using deformable templates. *Computer Vision and Image Understanding*, 5(11):303–319, 2000.

[14] Y. Shirai M. Etoh and M. Asada. Contour extraction by mixture density description obtained from region clustering. In *Proceedings of European Conference on Computer Vision*, pages 24–32, 1992.

[15] J. L. Prince and C. Xu. A new external force model for snakes. In *IEEE Proc. Conf. on Comp. Vis. Patt. Recog. (CVPR'97)*, pages 30–31, 2001.

[16] J. A. Sethian. Level set methods and fast marching methods. *Cambridge University Press*, 1999.

[17] A.L. Yuille. Cccp algorithms to minimize the bethe and kikuchi free energies: Convergent alternatives to belief propagation. *Neural Computation*, 14(7):1691–1722, 2004.

Oliviu Matei

Received October 11, 2006

Technical University of Cluj-Napoca,
E–mail: *oliviu.matei@holisun.com*