ISSN 1864-5909, Volume 3, Number 1

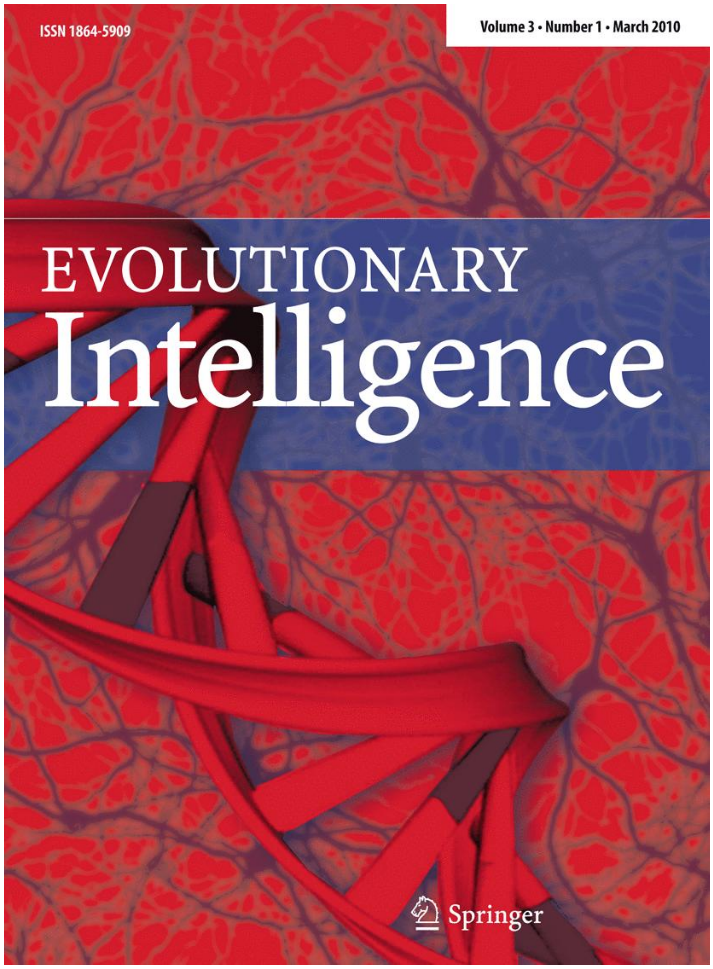

Evol. Intel. (2010) 3:1–11
DOI 10.1007/s12065-010-0036-x

RESEARCH PAPER

LSGA: combining level-sets and genetic algorithms for segmentation

Payel Ghosh · Melanie Mitchell · Judith Gold

Received: 10 June 2009 / Accepted: 5 February 2010 / Published online: 21 February 2010

Abstract A novel technique is presented to combine genetic algorithms (GAs) with level-set functions to segment objects with known shapes and variabilities on images. The individuals of the GA, also known as *chromosomes* consist of a sequence of parameters of a level-set function. Each chromosome represents a unique segmenting contour. An initial population of segmenting contours is generated based on the learned variation of the level-set parameters from training images. Each segmenting contour (an individual) is evaluated for its fitness based on the texture of the region it encloses. The fittest individuals are allowed to propagate to future generations of the GA run using selection, crossover and mutation. The GA thus provides a framework for combining texture and shape features for segmentation. Level-set-based segmentation methods typically perform gradient descent minimization on an energy function to deform a segmenting contour. The computational complexity of computing derivatives increases as the number of terms increases in the energy function. In contrast, here the level-set-based curve evolution/deformation is performed derivative-free using a genetic algorithm. The algorithm has been tested for segmenting thermographic images of hands and for segmenting the prostate in pelvic CT and MRI images. In this paper we describe the former; the latter is described in [11, 12]. The LSGA successfully segments entire hands on images in which hands are only partially visible. At the end of the paper we report experimental evaluation of the performance of LSGA and compare it with algorithms using single features: the Gabor wavelet based textural segmentation method [1, 9], and the level-set based segmentation algorithm of Chan and Vese [6].

Keywords Medical image processing · Image segmentation · Genetic algorithms · Level-set methods · Image texture analysis

P. Ghosh (✉) · M. Mitchell
Department of ECE, Portland State University, Portland, OR, USA
e-mail: payel@pdx.edu; payel@ece.pdx.edu

M. Mitchell
e-mail: mm@cs.pdx.edu

M. Mitchell
Department of Computer Science, Portland State University, Portland, OR, USA

M. Mitchell
Santa Fe Institute, Santa Fe, NM, USA

J. Gold
Department of Public Health, Temple University, Philadelphia, PA, USA
e-mail: jgold@temple.edu

1 Introduction

Genetic algorithms (GAs) [22, 32] simulate the learning process of biological evolution using selection, crossover and mutation. Genetic algorithms are blind optimization techniques that do not need derivatives to explore the search space. Instead they use payoff values, known as fitness, to guide the search towards better landscapes. This quality makes GAs more robust [15] than other local search procedures such as gradient descent or greedy techniques like combinatorial optimization.

GAs have been used for a variety of image processing applications, such as edge detection [18], image segmentation [19], image compression [34], feature extraction from remotely sensed images [19] and medical feature extraction [20]. The image processing problem being explored in this paper is image segmentation; a technique of delineating a region of interest on an image.

Level-set methods have become very popular in the field of image segmentation due to their ability to represent boundaries of objects that change with time or are ill-defined [35, 39]. In this method, a deformable segmenting curve is associated with an energy function and curve evolution/movement is performed by minimizing this energy function using gradient descent. The energy function may consist of region-based terms (such as pixel intensity values, edges etc.) and contour-based terms (such as curvature and length of the curve). These features are called low-level because they encode the information that can be derived directly from the image. There are many real-world problems that require high-level features for segmentation, such as the prior knowledge of shapes and context information derived from the extrapolations of human perception [4, 16]. Incorporating such information into an explicit energy function term may be difficult or impossible to encode for performing segmentation. A genetic algorithm (GA) solves this difficulty because it eliminates the energy function (and instead uses a fitness function) thereby providing a framework for incorporating high-level features and combining multiple features for segmentation. The level-set based genetic algorithm scheme (LSGA) proposed here uses the learned shape and textural properties of a known object to segment it on unseen images.

Here, the LSGA has been used to segment thermographic images of hands. The images of hands were acquired for studying upper extremity musculoskeletal disorders (UEMSD). The pathophysiology in UEMSDs is largely unknown. However, a component may include reduced blood flow in the upper extremity [3, 13, 14, 26, 36, 38]. Infrared thermography reveals skin temperature which is largely determined by subcutaneous perfusion. For this study subjects with UEMSD were given a typing challenge and images were taken at periodic time intervals. This segmentation problem is challenging because the fingers of the patients start to disappear on the thermographic images as their fingers become cold after typing (a symptom of UEMSD). A successful segmentation of these images involves deriving the prior human knowledge of the shape of the hand and modeling the movement of the hand and fingers from training images to perform the segmentation.

An individual in an LSGA population is a vector of parameters of a level set function and is referred to as a *chromosome* of the GA. The GA adapts the parameters of the level-set function to produce fit individuals or good segmentations of the given image using the information encoded in its fitness function. The algorithm terminates by finding a reasonable segmentation within the bounds of known shape and texture of the hands. Thus, the main contribution of this work is the use of GAs to optimize a level-set function thereby combining so-called high-level features (such as shape and texture) for segmentation.

The rest of the paper is organized as follows: first a literature review is provided on image segmentation specifically emphasizing level-set-based segmentation algorithms. The LSGA algorithm is then described in detail followed by a comparison with the level-set-based segmentation method of Chan and Vese [6] and with the Gabor wavelet-based segmentation method [1]. The description of the dataset used and the results achieved from applying the algorithm to thermographic images of the hand are then discussed. Discussion and evaluation of results is presented at the end.

2 Segmentation methods

Segmentation is defined as the process of demarcating an object on an image with a boundary/contour. Segmentation is performed by determining either pixel-level or object-level properties of an object that set it apart from the rest of the image. These properties can be edges, texture, pixel intensity variation inside the object, shape, size, orientation, and location of objects with respect to other objects in the image, and so on.

Pixel-based methods identify local features such as edges and texture in order to extract regions of interest on images. The most commonly-used pixel based operation is the edge-detector. Edges are defined as regions on the image with large pixel intensity variations. A comprehensive review of edge-detection methods is provided by [17]. However, these techniques can produce broken edges and also include boundaries of other features present in an image. Another intensity-based method is the region-growing method [42], which starts from a seed-point (usually placed manually) on the image and performs segmentation by clustering neighborhood pixels using a similarity criterion.

More complex pixel-level features are textures. Texture is usually defined as a region consisting of mutually related elements. Various approaches for textural feature extraction exist including co-occurrence matrices [10], filtering methods such as Gabor filters [9], Fourier transform methods, texture element finding methods such as textons [24, 40], Laws' texture method [27], to name a few. One major drawback of all pixel-based segmentation algorithms is that regions outside the object can also be identified as being part of the object and there is no notion of shape of a region in these methods.

Segmentation using object level-features involves quantifying object characteristics such as shape, pose [41], and relative position with respect to objects as well as region-based properties of the object. This quantification process transforms the object into a series of feature vectors [7]. Deriving these transformations and combining them is the main challenge of object-based segmentation. Object-level features are first derived from manually segmented objects and then used to perform object-level segmentation on test images. Shapes are generally represented using contours, transforms, or regions. In this work, shape has been represented using an active contour, which can be easily deformed to represent a flexible boundary; pose has been incorporated into the shape representation. The following section provides an overview of active contour methods of segmentation.

2.1 Active-contour method and the level-set method of segmentation

Deformable contour models (also known as active-contour models) are shape-based procedures that minimize an energy function to perform segmentation. The energy function is typically a function of regional properties of the image such as edges, mean pixel intensity, and/or object-level features such as curvature of the object and size. In these methods the initial contour is usually placed randomly or manually inside, on, or outside the region of interest. During the curve evolution process, minimization of the energy function drives the curve towards the boundary of the object.

One approach for curve evolution is the marker point method [35, 39] in which the segmenting curve C is parameterized by converting each point on the curve to represent a position vector $[s, t]$, where s are points of the curve along a certain orientation (clockwise or counter-clockwise), and t is time. The front can be interpolated from these marker points as either line segments in two-dimensions (2D) or triangles in three-dimensions (3D). One disadvantage of this method is that if the curve evolution makes two marker points come closer to each other into a corner then it can lead to an uneven advancement of the markers. Within a few time steps this can lead to oscillations in the curvature making the output unbounded.

Another approach to active shape modeling is the level-set method introduced by Sethian [39]. In this approach the evolving boundary (interface) is implicitly embedded as the zero isocontour of some function. For example, a unit circle can be defined as the zero isocontour of the function, $\varphi(\vec{x}) = x^2 + y^2 - 1$, $\varphi(\vec{x}) = 0$. In the level-set method, the equation of motion of the curve is defined using a simple convection equation (the level-set equation) such as:

$$\dot{\varphi} + \vec{V} \cdot \nabla\varphi = 0, \tag{1}$$

where $\dot{\varphi}$ is the temporal partial derivative of the implicit function ϕ, $\vec{V} = \langle u, v, w\rangle$ is the velocity field (u, v, w are components of the velocity field in the x, y and z directions, respectively), and ∇ is the spatial gradient operator.

The level-set function ϕ may be defined in terms of the signed distance function. The signed distance function takes any point in the plane and returns the Euclidean distance between the pixel and the closest point on the interface. Pixels outside the interface have positive distance while the pixels inside have negative distance values assigned to them. The zero level-set is defined as the interface itself, i.e., the set of all points that are at height zero, or equivalently, whose distance to the interface is zero. Curve evolution in the level-set method is stable and small errors in approximation are not amplified with time.

Level-set methods have been used extensively for image segmentation [8, 31]. Some of the popular methods are by Leventon et al. [28], Tsai et al. [43] and Chan and Vese [6]. Leventon et al. introduced the concept of shape representation by principal component analysis (PCA) on signed distance functions. They also incorporated statistical shape priors into their geodesic active-contour model to generate maximum *a posteriori* estimates of pose and shape. They segmented synthetic as well as medical images using their method and compared level-set evolution with and without shape influence. Their segmentation results were within one or two voxels of manual segmentation. However, the initialization point was placed manually on the images.

Chan and Vese introduced a region-based energy function based on Mumford-Shah segmentation techniques [33] in order to detect features with diffuse boundaries. The limitation of their model as pointed out by them in the paper is that it could only detect objects by intensity average values. They also mention that other image features such as texture need to be combined with a level-set framework in future to perform more generalized Mumford-Shah segmentation. The LSGA developed here attempts to address this need.

Tsai et al. derived a shape-based level-set function. Tsai et al.'s goal was to find the parameters of this function that produce a good model of the object shape based on priors from the training data. Tsai et al. derived these parameters via an optimization procedure that used statistics defined over local regions in a set of training images. The performance of Tsai et al.'s algorithm thus depended on the particular choice of statistics used to distinguish various regions within a given image. They showed automatic segmentation results on several synthetic images and semi-automatic segmentation on cardiac and pelvic MRI images. We have adopted this level-set function for the LSGA because it is a scheme for modeling known shapes as is required by the current problem. Here, a GA has been used to evolve the same level-set function using texture feature for curve evolution. In future, other high-level image features such as spatial relationships between objects would be tried to be incorporated into the LSGA.

2.2 Genetic algorithms for segmentation

Genetic algorithms have been used for segmentation by [2, 21, 29, 30, 37]. In [19, 20] a general-purpose image-segmentation system called GENIE (short for GENetic Imagery Exploration) is described that was developed at Los Alamos National Laboratory. GENIE used genetic programming to evolve image-processing "pipelines": sequences of elementary image processing operations, including morphological, arithmetic and point operators, filters and edge detectors, among others. Each pipeline, when run on a given multi-spectral image, performed image segmentation by classifying certain pixels as being part of a desired feature or otherwise. The fitness of each pipeline in the population was computed by comparing the final classification output with a set of training images, in which positive and negative examples of the desired feature had been manually highlighted. At the end of a run of GENIE, the fittest pipeline in the population was used to segment the desired feature in new images. A fitness measure similar to that used in GENIE has been adopted in the LSGA developed here because it has distinct payoff terms for reward and penalty. Harvey et al. [20] applied GENIE to a medical feature-extraction problem using multi-spectral histopathology images. Their specific aim was to identify cancerous cells on images of breast tissue. Their results were not very accurate, since GENIE used only texture-based image operations, and did not have any object or shape-based operators. Such operators are clearly needed for more accurate medical image segmentation.

A model-based image analysis technique using a GA is described in [29]. The method used a evolutionary Hough transform scheme to detect known shaped objects on images such as circle and ellipse. The GA population consisted of a set of points in the parameter space. In contrast, the LSGA evolves a population of segmenting contours constrained by known shape.

Cagnoni et al. [5] used a GA for segmenting images by evolving parameters of an active contour model called "snakes" [25]. It optimized an energy function based on low-level features such as smoothness of the curve, curvature and image gradient. In contrast, the LSGA framework evolves parameters of a level-set function. Unlike the explicit representation of shapes used in [5, 29], the level-set based implicit representation of shape used here allows textural features to be used for searching the parameter space.

3 LSGA: combining level-sets and genetic algorithms

The algorithm presented here consists of two stages: the training stage and the segmentation stage. In the training stage shape, shape variability and texture information of the region of interest are derived from manually segmented images. The data for the training stage is obtained from a set of training images on which a human has drawn a contour around the object to be segmented. The set of these training contours provides information about the shape and pose variability of the given object. The textural properties of the object are also derived from the same set of training data. The segmentation phase involves the genetic algorithm evaluating candidate contours for segmenting the desired object in a new image using a fitness measure, and iterating over successive generations until the fitness exceeds a threshold.

3.1 Training: deriving shape information

The shape representation is derived from the mean and variance of all manually drawn contours in a training set [43]. The manually drawn contours from the training data are first converted into signed distance functions, ψ_i ($i = 1$ to n, is the number of training contours). The level-set function is a linear combination of the mean shape and weighted shape variances in the signed-distance domain. The mean shape is defined for n contours as:

$$\bar{\Phi} = \left(\frac{1}{n}\right) \sum_{i=1}^{n} \psi_i. \tag{2}$$

Mean offset functions are derived by subtracting the mean from the signed distance representations of the training contours $\left(\tilde{\psi}_i = \psi_i - \bar{\Phi}\right)$. The columns of the mean offset functions (size $N = N_1 \times N_2$ the same as the training images) are then successively stacked on top of one another to form one large column vector (β_i) of size $1 \times N$. A new matrix $\boldsymbol{S}$ (size $N \times n$), called the *shape variability matrix*, is formed from n such column vectors

$$S = [\beta_1, \beta_2, \ldots, \beta_n]. \tag{3}$$

The variance in shape is then computed by an eigenvalue decomposition on this shape variability matrix as,

$$\frac{1}{n} SS^T = U \Sigma U^T. \tag{4}$$

Here $\boldsymbol{U}$ is an $N \times n$ matrix whose columns represent n orthogonal modes of shape variation and Σ is a diagonal matrix of eigenvalues. By rearranging the columns of $\boldsymbol{U}$ to form an $N_1 \times N_2$ structure, the n different eigenshapes can be obtained $\{\Phi_1, \Phi_2, \ldots, \Phi_n\}$. The mean shape and shape variability derived from the training phase are used to define a level-set function that represents the segmenting curve,

$$\Phi[w] = \bar{\Phi} + \sum_{j=1}^{k} w_j \Phi_j. \tag{5}$$

Here $\boldsymbol{w}$ are the weights for linearly combining the k principal eigenshapes. By incorporating pose parameters into this level-set framework a new level-set function is obtained that can handle object shapes with different sizes and orientation. Pose is defined using an affine transform which is the product of three matrices, the translation matrix, the scaling matrix and the rotation matrix, respectively,

$$\begin{bmatrix} \tilde{x} \\ \tilde{y} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & a \\ 0 & 1 & b \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} h & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}. \tag{6}$$

Here, x and y are the pixel coordinates of the input image and $\tilde{x}$,$\tilde{y}$ are the pixel coordinates of the affine transformed image. Note that a homogeneous coordinate system is used here. Using this homogeneous coordinate system allows the translation operation in Eq. (6) to be represented with a matrix multiplication.

This new level-set function is defined as [43]

$$\Phi[w,p](x,y) = \bar{\Phi}(\tilde{x},\tilde{y}) + \sum_{j=1}^{k} w_j \Phi_j(\tilde{x},\tilde{y}). \tag{7}$$

Here $\boldsymbol{p} = [a, b, h, \theta]$, a, b are x, y translation parameters, h is the scale factor and θ is the angle of rotation. The zero-level of this level-set function gives the segmenting contour and its parameters are evolved by the GA.

Before deriving the mean shapes and shape variance from the training data the images need to be aligned for pose. Gradient descent is used to minimize the difference between pairs of binary images with respect to their pose parameters. The transformed image based on pose is given by:

$$\tilde{I} = T[p] * I, \tag{8}$$

where, $\boldsymbol{T}[\boldsymbol{p}]$ is the 2D transformation matrix of Eq. (6). The energy functional used to minimize the difference between two images is given by:

$$E_{align} = \frac{\iint_{\Omega} (I^i - I^j)^2}{\iint_{\Omega} (I^i + I^j)^2}, \; i \neq j. \tag{9}$$

Here, Ω is the image domain. The area normalizing term in the denominator is employed to prevent the images from shrinking to improve the cost function. The initial pose parameters of one of the shapes are kept fixed and the pose parameters of the second image are calculated to minimize the pose differences.

3.2 Training: deriving texture information

The textural priors were derived from training images using Gabor wavelet-based texture segmentation method. Gabor wavelets are based on the Gabor elementary function given by the modulation of the Gaussian with a complex exponential function (Eqs. 10, 11).

$$h(x,y) = g(x,y)\exp\left[j2\pi\left(U_x + V_y\right)\right] \tag{10}$$

$$g(x,y) = \frac{1}{2\pi\sigma_x\sigma_y}\exp\left\{-\frac{1}{2}\left[\left(\frac{x}{\sigma_x}\right)^2 + \left(\frac{y}{\sigma_y}\right)^2\right]\right\}. \tag{11}$$

Gabor wavelets are derived from the mother wavelet $h(x, y)$ by several translations and dilations. The method of Gabor wavelets assumes that local texture regions are spatially homogeneous and the mean and standard deviation of the transform coefficients are used to represent regions for classification. The Gabor wavelet is given by Eq. (12)

$$\begin{aligned} h_{mn}(x,y) &= aH(x',y'), \\ a &> 1, \\ m,n &= \text{integers}, \\ \text{and } x' &= a^{-m}\left(xCos\frac{n\pi}{k} + ySin\frac{n\pi}{k}\right), \\ y' &= a^{-m}\left(-xSin\frac{n\pi}{k} + yCos\frac{n\pi}{k}\right). \end{aligned} \tag{12}$$

Here, k is the number of orientations. Given an image $I(x, y)$, the Gabor wavelet transform is given by

$$W_{mn}(x,y) = \int I(x_1,y_1)h^*_{mn}(x - x_1, y - y_1)dx_1dy_1. \tag{13}$$

Here, h^* is the complex conjugate of h. The mean and standard deviation of transform coefficients for each known region are the derived texture priors.

3.3 Segmentation using LSGA

Segmentation is performed using LSGA by optimizing a population of segmenting contours for shape, texture of enclosed region, location and pose. Each individual in the GA population is defined as a fixed-length string of real-valued *genes*.

$$I = [w_1, w_2, w_3, w_4, a, b, h, \theta]. \tag{14}$$

The four weight parameters are used for deriving the weighted $\pm\sigma_1$ and $\pm\sigma_2$ variation (where σ^2 is the eigenvalue corresponding to the principal eigenshapes) of the mean shape and a, b, h, and θ are pose parameters as defined in Eq. (6). For the individuals of the GA population, the pose parameters are chosen randomly from the space of real numbers. Each individual ($\boldsymbol{I}$) in the population represents a unique segmenting contour. This segmenting contour produces a binary image ($\boldsymbol{B}_1$) at its output with ones inside and zeros outside the segmenting contour.

The fitness of an individual $\boldsymbol{I}$ is calculated based on the degree to which the contour encloses the region of interest (ROI) in a test image (a new image not in the training set). The ROI on the test image is determined by the textural segmentation of the test image. Thus, LSGA performs deformable template matching around the textural region of interest. The fitness is calculated by comparing the two binary images $\boldsymbol{B}_1$ and $\boldsymbol{B}_2$. Binary ("true"/"false") image $\boldsymbol{B}_1$ is generated by the textural classification of pixels on a new test image using the mean and variance of the saved wavelet-coefficients. Each pixel of the test image is classified as "true" (ROI) or "false" (does not belong to ROI). The second binary ("true"/"false") image ($\boldsymbol{B}_2$) is obtained from the GA individual being evaluated, by placing the corresponding contour on the test image and classifying all pixels inside the contour as "true" and all other pixels as "false".

The fitness is a function of the detection rate (D) and the false alarm rate (F) as:

$$\text{Fitness} = 500(D + (1 - F)). \tag{15}$$

The detection rate is defined as the fraction of "true" ("false") pixels in the segmented image ($\boldsymbol{B}_2$) matching the "true" ("false") pixels in the textural classification $\boldsymbol{B}_1$. Note that $\boldsymbol{B}_1$ is the so-called "truth plane" used by the fitness function and is not the ground truth (derived from manual segmentation). The ground truth images are only used for evaluating the final segmentation results. The false alarm rate denotes the fraction of "false" ("true") pixels in $\boldsymbol{B}_1$ that are classified as "true" ("false") pixels in the segmented image $\boldsymbol{B}_2$. For convenience in calculations, the constant 500 scales the fitness so that the maximum fitness score that can be attained using this function is 1,000. The processes in GA evolution: selection, crossover, and mutation to create a new generation are iterated until the maximum fitness is attained or after a specified number of generations have been produced.

Rank selection and fixed-length crossover have been implemented here. Rank selection is implemented by comparing the fitness of individuals, and making individuals with higher fitness more likely to be selected to produce offsprings. Fixed-length crossover is performed by swapping fixed length segments of genes between two individuals. Mutation is performed by randomly changing the value of a gene with another real number within a fixed range of values. Mutation rate is defined as the probability of a single gene to be mutated. Similarly, the crossover rate defines the probability of a crossover to occur between two individuals.

For evaluating the performance of the algorithm the definitions of closeness of the segmentation outcome to the truth (here, manual segmentation) were derived using the dice similarity coefficient [44] and the partial Hausdorff distance [23]. Both of these measures have been extensively used for evaluating segmentation algorithms. The ground truth was obtained by averaging over multiple manual delineations.

The dice similarity coefficient provides a measure of the degree of overlap between two segmentations as:

$$DSC(A,B) = \frac{2|A \cap B|}{(|A| + |B|)}. \tag{16}$$

A DSC of 1 indicates a perfect match and 0 indicates no match. The partial Hausdorff distance is derived between the boundary points of two contours. If $\boldsymbol{A} = \{a_1,\ldots, a_p\}$ and $\boldsymbol{B} = \{b_1,\ldots,b_q\}$ be finite sets of points on two contours, then the partial Hausdorff distance between them is defined as:

$$H(A,B) = \max(h(A,B), h(B,A)), \tag{17}$$

where

$$h(A,B) = \max_{a \in A} \min_{b \in B} \|a - b\|. \tag{18}$$

The function $h(\boldsymbol{A},\boldsymbol{B})$ takes each point in $\boldsymbol{A}$ and finds the closest point in $\boldsymbol{B}$ from that point. It then ranks the points in $\boldsymbol{A}$ based on the distance values and finds the point with the greatest "mismatch". Thus, the partial Hausdorff distance is a measure of the distance by which two contours i.e., the final segmentation outcome and the ground truth differ.

4 Data: thermographic images

The data for this analysis has been obtained from Temple University's Ergonomics and Work Physiology Lab where researchers are studying musculoskeletal disorders of distal upper extremity (e.g., tendinitis and carpal tunnel syndrome). Infrared thermography of the hand reveals skin temperature which is related to the amount of blood flowing into the hand. For this study far-infrared images of hands were acquired using ThermaCAM AM40 thermographic camera (FLIR Systems, Wilsonville, OR) with a sampling rate of 7 Hz. The subjects were given a 9 min typing challenge and images were taken at the following intervals: before typing, 0–2 min (post-typing), 3–5 min (post-typing), and 8–10 min (post-typing).

Figure 1 (upper panel) shows the hands of four subjects before starting to type. Figure 1 (lower panel) shows the hands of the same patients 8 min after typing. The subjects'

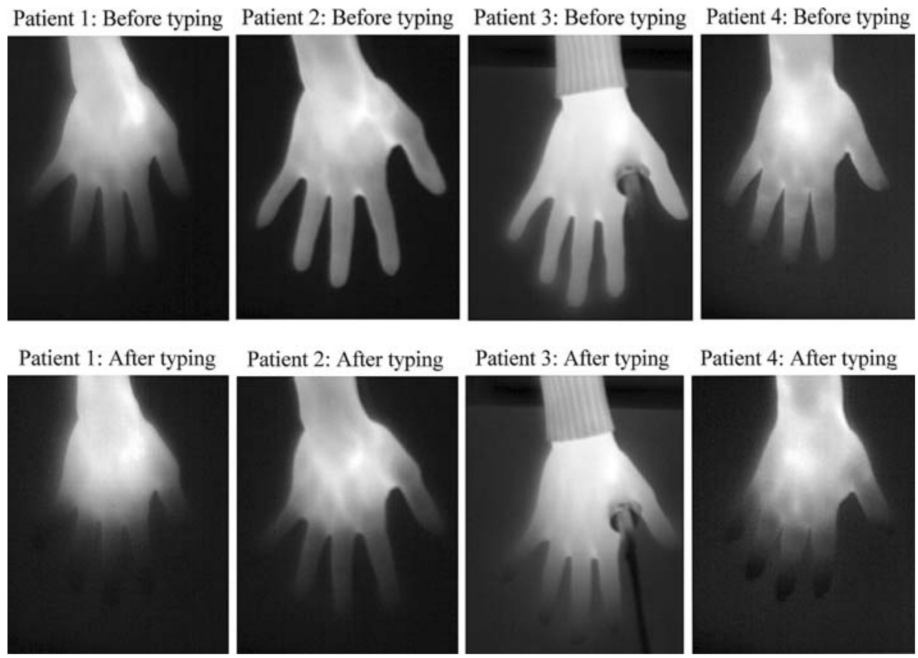

Fig. 1 Thermographic images of hands of four subjects taken before typing (*upper panel*); after 8 min of typing (*lower panel*). The fingers start to become invisible due to reduced blood flowing in the subjects' hands

fingers are partially visible on these images because the hands approach the temperature of the surrounding surface. Also a near infrared spectroscopy (NIRS) probe is adhered to the skin above the first dorsal interosseous muscle of their right hand (between the thumb and index finger) and on some of the images the boundary of the hand touching the probe is not visible.

These images were manually segmented by a human who has a prior knowledge of the shape of the hands and has also looked at the hand images of each subject taken prior to typing. The challenge in the problem is more than mere shape matching because the subjects tend to move their hands and fingers during the imaging process. Therefore, a rigid template matching method is not suitable for solving this problem. The LSGA performs a deformable template matching within known bounds of mean shape and shape variability (movement of fingers and the hand itself) and the texture of the region it encloses to perform the segmentation task.

Once an image is manually segmented a score of the mean temperature of the hand is generated and compared with the temperature of the hand at successive time steps to determine if and where the blood flow is changing in the hand with time. The dataset consists of images from four subjects. For each subject there are 500 images for each time range, that is, a total of about 2,000 images per subject.

5 Results and discussion

The experiments for segmenting the thermographic images of hands were set up in the following way. The images acquired prior to typing were used as the training images. The images taken after the typing challenge at various intervals were used as the test images. A subset of 240 test images (every 25th image was chosen i.e., about $20 \times 3 = 60$ images per patient) were used for validation purposes. The ground truth in the form of manual segmentation was derived only for the validation set. The segmentation performance for rest of the test images were analyzed visually by a human. A subset of the training images (~ 100) was also manually segmented by a human to derive the model for known shape, texture and movement of the hands of the subjects.

Figure 2 shows the mean shape and the variability of the mean shape from one patient. The eigenvalue σ_1 depicts the movement of the fingers and σ_2 the width of the fingers which varies between multiple segmentations (shown in Table 1 for all patients). Other eigenvalues affect the length of the fingers, the size of fingertips etc., and are ignored for this study because they are not the principal modes of variation of the shape of hands.

Segmentation was performed on the test images of each subject using the following methods: LSGA, Gabor wavelet based segmentation algorithm (GW), and the Chan and Vese algorithm (CV). Figure 3 (top left panel) shows a manually segmented hand test image. The outcome of the Gabor wavelet-based segmentation algorithm on the sample test image is shown in Fig. 3 top right panel. This method finds only the region of the hand visible by pixel intensity variations on the image. Figure 3 (bottom left panel) shows the result of applying the Chan and Vese level-set-based algorithm to the same hand image. Since both these methods do not have the notion of a known

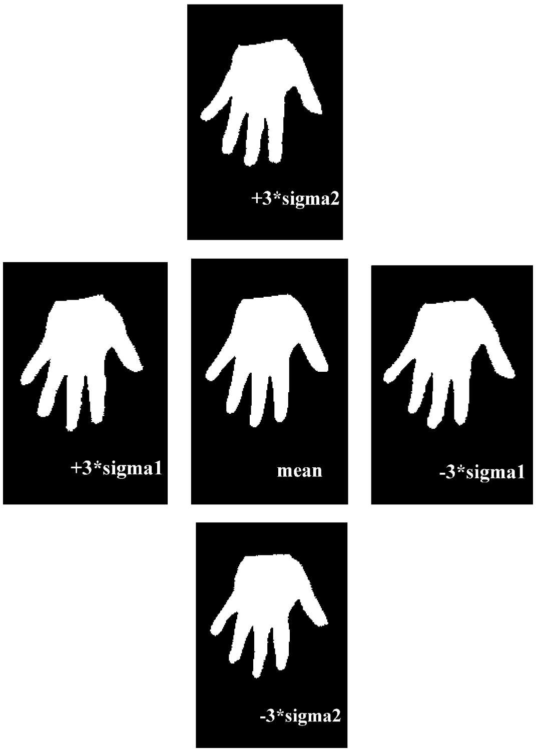

Fig. 2 Mean shape (*center*). $\pm 3\sigma_1$ variability depicting movement of fingers (*left*, *right*). $\pm 3\sigma_2$ variability of width of fingers (*top*, *bottom*)

Table 1 Shape variability of hands

	σ_1	σ_2
Patient 1	2.7×10^4	1.5×10^4
Patient 2	1.0×10^2	0.7×10^2
Patient 3	1.8×10^2	1.1×10^2
Patient 4	2.2×10^2	1.0×10^2

shape they could not segment the entire hand on the test image.

Finally, the LSGA was used to segment the same hand image. The parameters used by the GA to evolve the segmenting contour are shown in Table 2.

Figure 4 shows some candidate segmenting contours in the GA population. The LSGA found the optimal location and pose of the hand in the image from a population of 50 segmenting contours (Fig. 3 lower right panel). The segmenting contour is shown on top of the test image to show the segmentation outcome here. The fitness of the final segmenting contour for this image was 828. The figure clearly shows that the LSGA outperforms the other methods. This is also confirmed by visual analysis of the results.

Figures 5, 6, and 7 show how the fitness values change as the GA parameters such as mutation rate, crossover-rate and population size vary. Therefore, a population size of 50, mutation rate of 10% and a crossover rate of 50% were

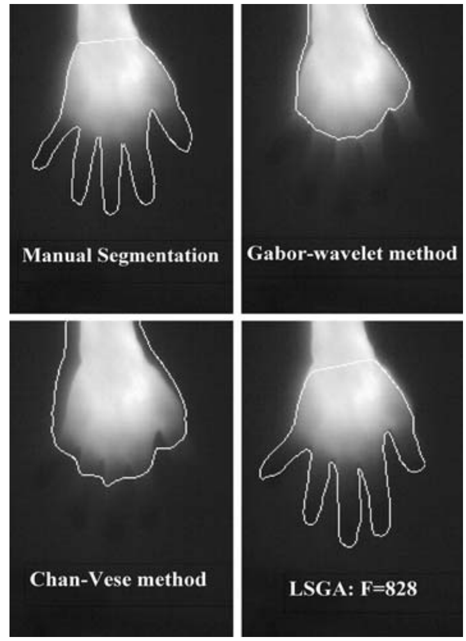

Fig. 3 Segmentation outcome on a test image (manual segmentation shown in *upper left*) using: Gabor-wavelet based method (*upper right*), Level-set based method of Chan and Vese (*lower left*), LSGA (*lower right*)

Table 2 LSGA parameters

Population size	50
Mutation rate	10% per gene
Crossover rate	50% single-point
Selection criteria	Rank selection
Weights for eigen shapes, *w*	(1–5) integers
Translation parameters *a*, *b*	Integer (0–30)
Rotation parameter θ	$-90°$ to $+90°$
Scale parameter *h*	1

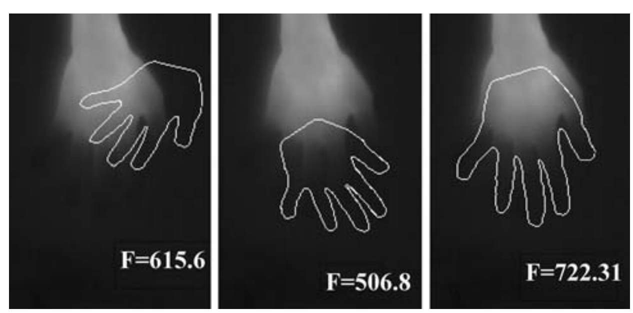

Fig. 4 Three candidate segmenting contours in the GA population. Here, *F* denotes the fitness of each individual defined in Eq. (15)

chosen as the default parameters for performing segmentation on all the images. Figure 8 shows the segmentation results from every time interval for each patient.

The average *DSC* and *H* were computed from all the segmentation outcomes for each patient. Table 3 shows

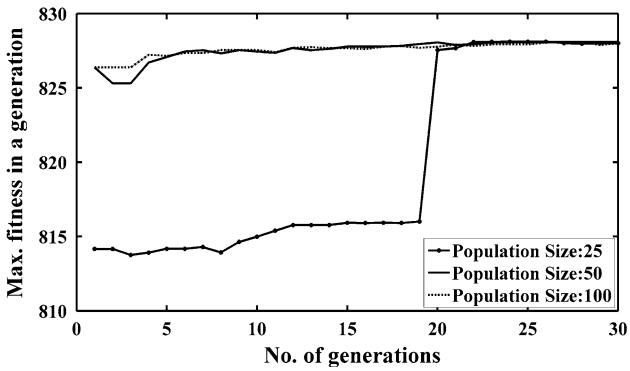

Fig. 5 Variation of the maximum fitness by number of generations of the GA run for the population sizes of 25, 50, and 100

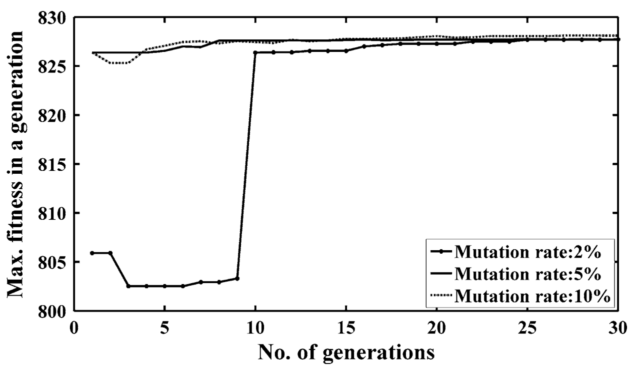

Fig. 6 Variation of the maximum fitness by number of generations of the GA run for the mutation rates of 2, 5, and 10%

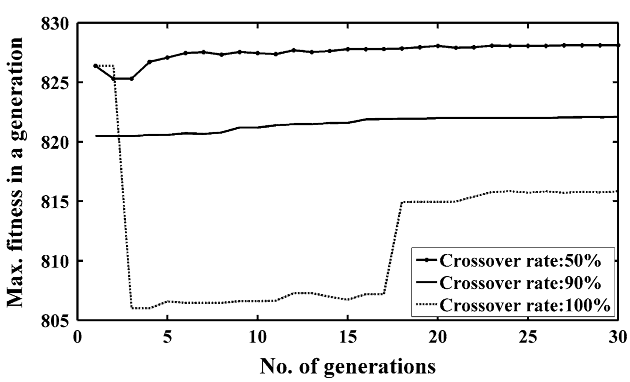

Fig. 7 Variation of the maximum fitness by number of generations of the GA run for the crossover rates of 50, 90, and 100%

the average values of *DSC* and *H* obtained for all the images in the validation set for each of the four subjects. Note that in the case of patient 4 the *DSC* is low and for patient 1 the *H* is relatively high even though a visual inspection of the results shows satisfactory segmentation. This is due to the limitation in acquiring accurate ground truth as the fingertips are completely invisible in the test images.

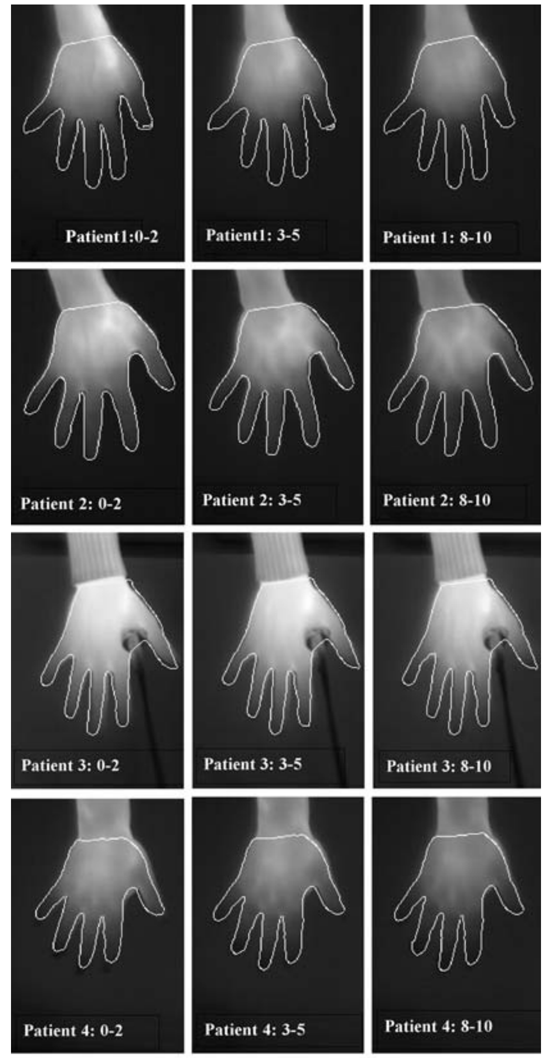

Fig. 8 Segmentation result from every time interval for each patient

Table 3 Validation criteria: comparison with ground truth

	DSC	*H*
Patient 1	0.8	6.7
Patient 2	0.9	3.4
Patient 3	0.85	3.5
Patient 4	0.6	3.5

6 Conclusion and future work

The LSGA performs derivative-free optimization of a level-set function for image segmentation. This brings flexibility to the level-set curve evolution process by letting the user choose different kinds of features for exploring the fitness landscape. In this paper two types of features, texture and shape, have been explored for evolving the level sets for segmenting thermographic images of hands.

Although these images had visible textural areas separating the hand region from the background, the knowledge of known shape was needed to segment them. The LSGA combined shape with texture to achieve the desired segmentation.

The current LSGA framework has been used to segment the prostate gland in pelvic CT/MRI images, a more complex segmentation problem, with promising results [11, 12]. However, features such as relative positions of objects need to be incorporated into this framework in the future to produce more accurate segmentations. This paper presents the basic framework for an evolutionary image segmentation scheme using level-sets which can be improved further for addressing more complex segmentation tasks.

References

1. Ahmadian A, Mostafa A (2003) An efficient texture classification algorithm using Gabor wavelet. In: Proceedings of the 25th annual international conference of the IEEE EMBS Cancun, Mexico, 17–21 Sept, pp 930–933
2. Ballerini L (1999) Genetic snakes for medical images segmentation. In: Poli R et al. (eds) Proceedings of the first European workshops on evolutionary image analysis, signal processing and telecommunications, Lecture notes in computer science, 1596. Springer-Verlag, London, 26–27 May 1999, pp 59–73
3. Brunnekreef JJ, Oosterhof J, Thijssen DHJ, Colier WN, Van Uden CJ (2006) Forearm blood flow and oxygen consumption in patients with bilateral repetitive strain injury measured by near-infrared spectroscopy. Clin Physiol Funct Imaging 26:178–184
4. Bunke H (2000) Graph matching for visual object recognition. Spat Vis 13:335–340
5. Cagnoni A et al (1999) Genetic algorithm-based interactive segmentation of 3D medical images. Image Vis Comput 17(12):881–895
6. Chan T, Vese L (2001) Active contours without edges. IEEE Trans Image Proc 10:266–277
7. Costa LF, Cesar RM Jr (2001) Shape analysis and classification theory and practice. CRC Press, Boca Raton
8. Cremers D, Rousson M, Deriche R (2007) A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. Intl J Comp Vis 72(2):195–215
9. Dunn D, Higgins WE (1995) Optimal Gabor filters for texture segmentation. IEEE Trans Image Process 4(7):947–964
10. Elfadel IM, Picard RW (1993) Gibbs random fields, co-occurences and texture modeling. IEEE Trans Patt Anal Mach Intell 16(1):24–37
11. Ghosh P, Mitchell M (2006) Segmentation of medical images using a genetic algorithm. In: Proceedings of the 8th annual conference on genetic and evolutionary computation, (Seattle, Washington, USA, 08–12 July 2006). GECCO '06, ACM, New York, pp 1171–1178
12. Ghosh P, Mitchell M, Tanyi J, Hung A (2009) A genetic algorithm-based level-set curve evolution for prostate segmentation on Pelvic CT and MRI Images. In: Gonzalez F, Romero Eduardo R (eds) Biomedical image analysis and machine learning, applications and techniques, chap. 6. IGI Global, Hershey
13. Gold JE, Cherniack M, Hanlon A, Dennerlein JT, Dropkin J (2009) Skin temperature in the dorsal hand of office workers and severity of upper extremity musculoskeletal disorders. Int Arch Occup Environ Health (epub)
14. Gold JE, Cherniack M, Buchholz B (2004) Infrared thermography for examination of skin temperature in the dorsal hand of office workers. Eur J Appl Physiol 93(1–2):245–251
15. Goldberg DE (1989) Genetic algorithms in search, optimization and machine learning. 1st Addison-Wesley Longman Publishing Co Inc
16. Goldberger J, Greenspan H (2006) Context-based segmentation of image sequences. IEEE Trans Patt Ana Mach Intell 28(3):463–468
17. Gonzales RC, Woods RE (2002) Digital image processing, 2nd edn. Prentice Hall
18. Harris C, Buxton B (1996) Evolving edge detectors. Research note RN/96/3, University College London, Dept. of Computer Science, London
19. Harvey N et al (2002) Comparison of GENIE and conventional supervised classifiers for multispectral image feature extraction. IEEE Trans Geosci Remote Sens 40(2):393–404
20. Harvey N, Levenson RM, Rimm DL (2003) Investigation of automated feature extraction techniques for applications in cancer detection from multi-spectral histopathology images. Proc SPIE 5032:557–566
21. Hill A, Taylor C (1992) Model-based image interpretation using genetic algorithms. Image Vis Comput 10(5):295–300
22. Holland JH (1975) Adaptation in natural and artificial systems. University of Michigan Press, Ann Arbor
23. Huttenlocher D, Klanderman G, Rucklidge W (1993) Comparing images using the hausdorff distance. IEEE Trans Patt Analysis Mach Intell 15:850–863
24. Julesz B (1986) Texton gradients: the texton theory revisited. Biol Cybern 54:245–251
25. Kass M, Witkin A, Terzopoulos D (1988) Snakes: active contour models. Intl J Comp Vis 1:321–331
26. Larsson R, Öberg PÅ, Larsson S-E (1999) Changes of trapezius muscle blood flow and electromyography in chronic neck pain due to trapezius myalgia. Pain 79:45–50
27. Laws KI (1980) Texture image segmentation. PhD dissertation, University of Southern California
28. Leventon M, Grimson E, Faugeras O (2000) Statistical shape influence in geodesic active contours. Proc IEEE Conf Comp Vis Patt Recog 1:316–323
29. Louchet J (2007) Model-based image analysis using evolutionary computation. In: Genetic and evolutionary computation for image processing and analysis. Hindawi
30. MacEachern L, Manku T (1998) Genetic algorithms for active contour optimization. IEEE Proc Intl Symp Circuits Sys 4:229–232
31. Malladi R, Sethian JA, Vemuri BC (1995) Shape modeling with front propagation: a level set approach. IEEE Trans Patt Anal Machine Intell 17(2):158–175
32. Mitchell M (1996) An introduction to genetic algorithms. MIT Press, Cambridge
33. Mumford D, Shah J (1989) Optimal approximation by piecewise smooth functions and associated variational problems. Commun Pure Appl Math 42:577–685
34. Nordin P, Banzhaf W (1996) Programming compression of images and sound, in genetic programming. In: Koza JR et al (eds) Proceedings of the 1st annual conference. Morgan Kaufmann, San Francisco
35. Osher SJ, Fedkiw RP (2002) Level set methods and dynamic implicit surfaces, 1st edn. Springer, New York
36. Oskarsson E, Gustafsson B-E, Pettersson K, Piehl Aulin K (2007) Decreased intramuscular blood flow in patients with lateral epicondylitis. Scand J Med Sci Sports 17:211–215

37. Poli R, Cagoni S (1997) Genetic programming with user-driven selection: experiments on the evolution of algorithms for image enhancement. In: Koza JR et al. (eds) Genetic programming 1997, proceedings of the 2nd annual conference, Morgan Kaufmann, San Francisco
38. Pritchard MH, Pugh N, Wright I, Brownlee M (1999) A vascular basis for repetitive strain injury. Rheumatology 38:636–639
39. Sethian JA (1999) Level set methods and fast marching methods, 2nd edn. Cambridge University Press, Cambridge
40. Shotton J, Winn J, Rother C, Criminisi A (2006) Textonboost: joint appearance, shape and context modeling for multi-class object recognition and segmentation. In: European conference on computer vision. pp 1–15
41. Sonka M, Hlavac V, Boyle R (1994) Image processing, analysis and machine vision. Chapman and Hall, London
42. Tremeau A, Borel N (1997) A region growing and merging algorithm to color segmentation. Pattern Recogn 30(7):1191–1203
43. Tsai A et al (2003) A shape-based approach to the segmentation of medical imagery using level sets. IEEE Trans Med Imaging 22:137–154
44. Van Rijsbergen CJ (2004) The geometry of information retrieval. Cambridge University Press, Cambridge