(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
24 December 2008 (24.12.2008)

PCT

(10) International Publication Number
WO 2008/155255 A1

(51) International Patent Classification:
G06T 5/00 (2006.01)

(21) International Application Number:
PCT/EP2008/057188

(22) International Filing Date: 10 June 2008 (10.06.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
07110531.6 19 June 2007 (19.06.2007) EP

(71) Applicant *(for all designated States except US)*: **AGFA HEALTHCARE** [BE/BE]; IP Department 3802, Septestraat 27, B-2640 Mortsel (BE).

(72) Inventor; and

(75) Inventor/Applicant *(for US only)*: **DEWAELE, Piet** [—/BE]; Agfa-gevaert, IP Department 3802, Septestraat 27, B-2640 Mortsel (BE).

(74) Common Representative: **AGFA HEALTHCARE**; IP Department 3802, Septestraat 27, B-2640 Mortsel (BE).

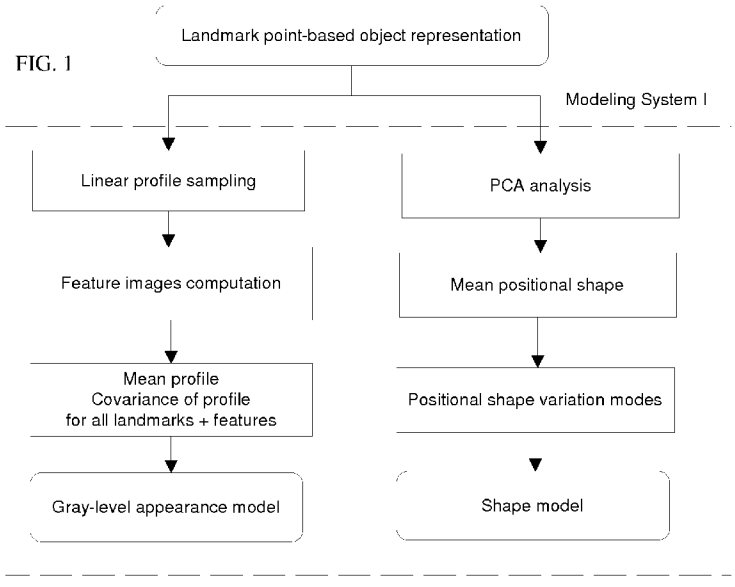
(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

(54) **Title:** METHOD OF CONSTRUCTING A GRAY VALUE MODEL AND/OR A GEOMETRIC MODEL OF AN ANATOMIC ENTITY IN A 3D DIGITAL MEDICAL IMAGE

FIG. 1



(57) **Abstract:** A gray value model is generated encoding photometric knowledge at landmark positions. This step exploits intensity correlation in neighborhoods sampled around landmark positions. A geometric model is generated encoding geometric knowledge between landmarks. This step exploits spatial correlation between landmarks of segmented anatomic entities. The topology of the landmarks on the surface or volume of the anatomic entity is encoded in a graph structure.

WO 2008/155255 A1

Method of constructing a gray value model and/or a geometric model of an anatomic entity in a 3D digital medical image.

[DESCRIPTION]

FIELD OF THE INVENTION

The present invention relates to a method to construct a photometric (also called gray value model) and/or geometric model (also called shape model) of anatomic entities from a training set of digital medical images.
Such models can be applied in a model-based segmentation procedure to segment the anatomic entities in a newly acquired image. Measurements of anatomy may be based on characteristic points of the segmentation.

BACKGROUND OF THE INVENTION

In radiological practice, geometric measurements are frequently used to aid diagnosis of abnormalities. In order to perform these measurements, key user points must be placed in an image, for example in an image displayed on a display device, on their corresponding anatomical landmark position. Measurements such as the distance between two points, or the angulation between lines are based on the position of these key user points. Furthermore, the geometry as a whole may be assessed for normality or abnormality, involving an analysis of the complete shape. Hence there is a need to automate and objectify the extraction of quantitative information that is embedded in a radiological image.

An example of such a frequently performed measurement is the computation of the cardiothoracic ratio (CTR) in thorax RX images (fig. 6). This ratio is defined as the ratio of the transverse diameter of the heart, at the level of the apex, to the internal diameter of the thorax (ID), i.e. CTR=(MLD+MRD)/ID.

The transverse diameter of the heart is composed of the maximum transverse diameter on the left side of the heart (MLD) and the maximum transverse diameter on the right side of the heart (MRD). Clearly, this definition entails that the radiologist searches along the inner border of the left and right ribcage boundary to locate the point pair needed to compute the internal diameter ID. This point pair must lie at the greatest internal diameter of the thorax. Likewise, the left and right heart shadow border must be searched to locate the points needed to compute the sum of MLD and MRD. More specifically, these points are situated most distant with respect to the midline of the spine. The process of border search requires that the radiologist is performing anatomy segmentation and locating the points (a total of four in this example) on the segmented anatomy. The segmentation step, in the case of CTR computation, amounts to delineating the lung fields.

An example of a 3D application is the segmentation of the lungs and the computation of the lung volume.

Many other measurements in digital images follow a similar approach involving the segmentation of the anatomic organ or entity onto which segmented geometry the characteristic points and measurement objects are determined and finally measurements are performed (fig. 5).

Referring to the example of cardiothoracic index calculation, to automatically position the required points, a method is needed to automatically segment the lung fields on a chest radiographic image.

The segmentation problem can be approached in several ways, depending on the application. Segmentation strategies evolved from low-level strategies in the early years of computer vision to the more recent model-based strategies.

Low-level methods rely on local image operators separating pixels with different photometric characteristics and grouping of pixels

with similar local photometric characteristics. Examples of both classes are edge detection and region growing. Despite the poor performance of these low-level approaches, they are very popular in most commercial image analysis tools. The main reasons are that they are simple to understand and to implement. For complex image data however, such as present in medical images and exemplified by the content of a thorax image as described above, their usefulness is limited.

More successful methods incorporate a priori knowledge about the shape to be segmented and about the photometric or gray-level appearance of the object in the image. These methods, referred to as model-based methods are often based on template matching. A template is matched for instance by correlation or with generalized Hough transform techniques. Unfortunately, the template matching is likely to fail in case of medical images. This is due to the large variability in shape and gray-level appearance that the anatomic object may exhibit.

Methods based on active contours, introduced by Kass et. al. (M. Kass, A. Witkin, and D. Terzopoulos, Snakes: active contour models, Int. J. Computer Vision, 1(4):321-331, 1988) and level sets (J.A. Sethian, Level set methods and fast marching methods, Cambridge Univ. Press, Cambridge, U.K. 1999) are able to cope with a larger shape variability, but are still unsuited for many medical segmentation tasks because little a priori knowledge about the object to be segmented can be incorporated.

Handcrafted parametric models overcome this problem, but are limited to a single application.

In view of these shortcomings, it is obvious that there is need for a generic segmentation scheme that can be trained with examples in order to acquire knowledge about the shape of the object to be segmented and the gray-level appearance of the object in the image. Active shape models (ASMs), introduced by Cootes and Taylor (T.F. Cootes, C.J. Taylor, D. Cooper, J. Graham, Active Shape Models - their training and applications, Computer Vision and Image

Understanding, 61(1):38-59, 1995) satisfy this definition of segmentation schemes. The shape model is given by the principal components of vectors of landmark points. The gray-level appearance model describes the statistics of the normalized first derivative of profiles centered at each landmark that run perpendicular to the object contour. The location of a landmark in a new image is found by minimizing the Mahalanobis distance between the first derivative profile and the distribution of the profile. This algorithm starts from an initial estimate and performs a fitting procedure, which is an alternation of landmark displacements and shape model fitting. Similar approaches have been devised all employing a three-step procedure. First, they all use a shape model that ensures that plausible results are generated. Secondly, they use a gray-level appearance model to place the object at a location where the gray-level pattern around the border or within the object is similar to what is expected from the training examples. Finally, the algorithm fits the model by minimizing some cost function.

The approach presented by active shape models is still faced with several limitations.

A first limitation is the need for an initial estimate, known as the initial positioning problem. In some cases, an extensive search for a proper initial contour is needed in this framework.

A second limitation lies in the alternate use of shape model and gray-level appearance model. Firstly, the segmentation estimate is updated using the gray-level appearance model. Secondly, the shape model is fitted to the updated contour. Unfortunately the shape model is misled if the gray-level appearance model wrongly locates a landmark.

Some experiments with active shape model segmentation schemes exhibited another problem. If the gray-level appearance model for a specific landmark is used to find a better location for that landmark, it requires that the true landmark location lies inside

the region the algorithm is exploring. In some cases the algorithm may search for a landmark in the region of another landmark. This means that the wrong gray-level appearance model is used at the wrong region, resulting in a poor segmentation.

It is an object of the present invention to provide an improved method of constructing a photometric and/or a geometric model of an anatomic entity (anatomic entities) in a digital medical image.

SUMMARY OF THE INVENTION

The above-mentioned aspects are realised by a method as set out in claim 1.

Specific features for preferred embodiments of the invention are set out in the dependent claims.

The gray value model comprises the probability distribution of a gray value and a number of multiscale derivatives of gray values at a set of landmarks distributed over an anatomic outline.
The shape model comprises the probability distribution of each connection vector between successive landmarks along the outline of an anatomic entity.

In the following the terms gray value model, gray value appearance model, gray level model, gray level appearance model and photometric model are used as synonyms.

Likewise the terms shape model and geometric model are used as synonyms.

The terms feature image and feature are also used as synonyms.

The embodiments of the methods of the present invention are generally implemented in the form of a computer program product

adapted to carry out the method steps of the present invention when run on a computer.
The computer program product is commonly stored in a computer readable carrier medium such as a DVD or a CD-ROM. Alternatively the computer program product takes the form of an electric signal and can be communicated to a user through electronic communication.

In EP A 1349098 a method is disclosed to automate the measurements in digitally acquired medical images by grouping measurement objects and entities into a computerized measurement scheme consisting of a bi-directionally linked external graphical model and an internal informatics model.
In a measurement session according to EP A 1349098, a measurement scheme is retrieved from the computer and activated. Measurements are subsequently performed on the displayed image under guidance of the activated measurement scheme.
In this computerized method, a multitude of geometric objects are mapped in the digital image onto which other measurement objects and finally measurement entities (such as distances and angles) are based. The basic geometric objects are typically key user points, which define other geometric objects onto which geometric measurements are based.
The above-identified patent application does not disclose however how the mapping may be effectuated automatically without the need of user positioning and manipulation of key measurement points. The models according to the present invention can be used in a method to effectuate this mapping.

The two-dimensional case has been described in the published European patent applications EP 1 760 658 and EP 1 760 659, where the segmentation results in one or more contours, open-ended or closed and enclosing the anatomical object of interest.

A discrete point-based object representation is introduced to describe the anatomic outline in a medical image.

Two strategies and associated systems for generating models and application of the models to segment actual images are disclosed. These strategies, although decomposing in building blocks with equal goals, construct and employ different photometric and geometric model knowledge.

First, a new gray level appearance model is constructed from training images, encoding photometric knowledge at landmark positions. This step exploits intensity correlation in neighborhoods sampled around each landmark.

Secondly, a shape model is constructed, encoding geometric knowledge between landmarks. This step exploits spatial correlation between landmarks of the segmented objects.

In a further segmentation step, photometric and geometric knowledge can jointly be used to segment one or more anatomic contours on a new image.

The resulting segmentations may be used to derive the position of desired measurement points in the course of performing geometric measurements on an image.

Measurement points can be defined in either a direct or an indirect way.

In the direct way, the positions of the landmarks are used as individual measurement points.

In the indirect way, the landmarks resulting from the segmentation process are interconnected by a curve and characteristic points on this curve are defined as measurement points.

Gray level appearance models

The first step in each modeling system constrains the geometric position of the landmark to lie on a linear path perpendicular to the present contour or a sparse grid centered on the landmark. The photometric pattern at a landmark point in an image is captured by a profile vector sampled in feature images based on image derivates extracted at a number of scales. The profile sampling may be along a linear path in the images or on a circular path around the landmark point.

Feature images

When the photometric pattern is computed for all image points in the image, feature images are obtained, the size of which is equal to that of the original image. Neighborhood operations are typically employed to compute the feature images. The values of the features in the feature image for a specific point can be arranged in a feature vector. The computation of feature images typically comprise two steps: (a) a linear filtering step using a bank of convolution kernels and (b) a post-processing step involving the sub-steps of a non-linear point operation combined with the computation of local image statistics. In step (b) one of the sub-steps may be performed alone.

Different filters and different types of post-processing lead to different types of feature images. For example, Laws constructed 25 two-dimensional convolution kernels as the outer product of pairs of vectors resembling ideal feature profiles such as Level, Edge, Spot, Wave and Ripple. The post-processing step consists of non-linear windowing in which the absolute filtered values are averaged in larger windows to obtain image texture energy measures. Unser used filter banks constructed as the outer product of vectors corresponding to well-known transforms like discrete sine (DST), discrete cosine (DCT), discrete even sine (DEST), discrete real even Fourier (DREFT), discrete real odd Fourier (DROFT) transforms, and for post-processing the computation of the channel variances. Finally, Gabor filter banks are constructed using symmetric or anti-symmetric Gabor filters tuned to combinations of different spatial frequencies and different orientations, such that all outputs have equal frequency bandwidth. The non-linear point operation step may consist of thresholding in this case.

Other image characteristics may be measured using feature images based on multi-resolution wavelets filters, rank-value filters or adaptive filters.

Feature images based on locally orderless images (LOI's)

The image structure around an anatomical landmark in a medical image exhibits a typical gray-level pattern. In the present invention, this structure is captured in a number N of mathematical features as outlined in the sequel. The concept of using feature images to build a gray-level appearance model was introduced earlier in B. Van Ginneken et al., Active shape model segmentation with optimal features, IEEE Trans. on Medical Imaging, 21(8):924-933, 2002, based on so-called locally orderless images (LOI's) presented by J.J. Koenderink and A.J. Vandoorn, The structure of locally orderless images, Int. J. of Computer Vision, 31(2):159-168, 1999.

A Taylor expansion can be used to approximate a gray-level function f around a point of interest, i.e. a landmark, at position x_0, by a polynomial of order K. The coefficients of each term are given by the derivatives $f^{(i)}$ at x_0:

$$f(x) \approx \sum_{i=0}^{K} \frac{1}{i!} f^{(n)}(x_0)(x - x_0)^i$$

All information in this function approximation is contained in the derivatives $f^{(i)}$. Similarly, the image can be described by an image filter bank containing the derivatives of the original image $\left(L, L_x, L_y, L_{xx}, L_{yy}, L_{xy}, \ldots\right)$. Each derivative image may further be computed at a number of blurring or inner scales σ. The feature images are finally obtained by computing the post-processing step as local image statistics of each derivate image. The local image statistics consist of a number of moments of the local histogram in a window with width α around each location x_0. In a specific embodiment the first and second histogram moments, all derivatives up to second-order $\left(L, L_x, L_y, L_{xx}, L_{yy}, L_{xy}\right)$ and 5 inner scales (σ=0.5,1,2,4,8 pixels) are computed, amounting to a total of $N = 2 \times 6 \times 5 = 60$ feature images. The window scale to compute the histogram is in accordance to the inner scale, i.e. $\alpha = 2\sigma$.

The limitation of using the derivatives in the row and column directly (i.e. x resp. y) in the feature vector is the lack of invariance with respect to translation and rotation (i.e. Euclidean Transformations). Unless the images are anti-rotated into a standard orientation, these operators can only be trained on examples that have the same orientation. Cartesian differential invariants (CDI) describe the differential structure of an image independently of the chosen Cartesian coordinate system; hence the anti-rotation step can be omitted. The construction of CDI's from Gaussian differential operators was described in L.M.J. Florack et al., Scale and the differential structure of images, Image and Vision Computing, Vol. 10 (6):376-388, 1992. The following are CDI's using derivatives up to second order: $\left(L, L_{xx}+L_{yy}, L_x^2+L_y^2, L_x^2 L_{xx}+2L_{xy}L_xL_y+L_y^2 L_{yy}, L_{xx}^2+2L_{xy}^2+L_{yy}^2\right)$. Again, at each pixel location, feature images are constructed using these operators at number of inner scales σ and computing the first two moments of the locally orderless histograms in windows of extent α, with α linked to the inner scale, e.g. $\alpha = 2\sigma$. These operators are independent to Euclidean transformations of image geometry, but still depend on the choice of the scale parameter σ.

Profile and feature similarity measurement

The gray level appearance model is built on the basis of a mean profile and variance-covariance matrix, and exploits intensity correlation between points of the profile, as explained next. The intensity at a profile is characterized by the intensity feature vector as outlined above.

In order to compare a profile with another profile, it is necessary to know how each point in the profile relates to any other point, and to measure the extent by which the value of a certain point of the profile varies from the mean with respect to any of the other points. This statistical relationship is captured by the covariance measure, which is always computed between 2 points (or the features at those points). If the covariance is measured between the feature value of a point and itself, the variance is obtained. When the

profile consists of k points at either side of the contour, for a total of $2k+1$ points, $(2k+1)k$ number of relationships exists. The covariance between a pair of points from the profile tells the way how a first value of the pair varies with respect to the second. When the covariance is large and positive, it indicates that when the first point's feature value increases, so does the second in a proportional way. When the covariance is large and negative, it indicates that when the first point's feature value increases, the second decreases in a proportional way. When the covariance is zero, it means that there is no systematic coupling in the feature values, i.e. they are independent of each other. All variance-covariance pairs can be arranged in a covariance matrix S, which is symmetrical about the main diagonal. The covariance matrix thus captures all structural dependencies between feature value pairs in the profile. In practice, the covariance matrix is built from learning samples, in this case a collection of profiles at a number of landmarks and a number of intensity features in a set of training images.

The covariance relates to how distance is calculated in higher dimensional space, i.c. the space of the $2k+1$ points in a given profile, to compute the distance between a current profile and a model profile. When the profiles would only comprise a single point with value x, the profiles could be compared with the model point value $\bar{x}$ for similarity just by subtracting the values of the current profile point and model profile point and taking the absolute value, i.e. $d=|x-\bar{x}|$. When the profiles would comprise two points, the distance between the 2-vector $\mathbf{x}$ and the model 2-vector $\bar{\mathbf{x}}$ could be computed as the Euclidean distance in the plane, i.e.

$$d=\sqrt{(\mathbf{x}-\bar{\mathbf{x}})'(\mathbf{x}-\bar{\mathbf{x}})}=\sqrt{|x_1-\bar{x}_1|^2+|x_2-\bar{x}_2|^2}\ .$$

Euclidean distance weighs all directions equally. However, when the variables x_1 and x_2 have unequal standard deviations σ_1 resp. σ_2, the contours of iso-distance lines are ellipses with major axes parallel to the coordinate axes. The differences may be weighted

with the inverse of the standard deviation, resulting in the weighted Euclidean distance

$$d_w = \sqrt{(\mathbf{x}-\bar{\mathbf{x}})' W (\mathbf{x}-\bar{\mathbf{x}})} = \sqrt{\left|\frac{x_1-\bar{x}_1}{\sigma_1}\right|^2 + \left|\frac{x_2-\bar{x}_2}{\sigma_2}\right|^2},$$

The matrix W contains the inverse of the variances of the variables, i.e.

$$W = \begin{bmatrix} 1/\sigma_1^2 \\ 1/\sigma_2^2 \end{bmatrix}.$$

Intensity levels of pixels close to each other as in the case of linear or circular profiles will be heavily correlated when they belong to an identical anatomical region with smoothly varying intensity levels. When pixel pairs are considered belonging to different tissues, the correlation may be inversely proportional. In either case, imaging noise and unimportant anatomical variation introduces a non-zero variance for the intensity level of the pixel itself.

Hence, if the variables are also correlated and have different variances, the inverse of the covariance matrix must be inserted, yielding the so-called Mahalanobis distance

$$d_m = \sqrt{(\mathbf{x}-\bar{\mathbf{x}})' S^{-1} (\mathbf{x}-\bar{\mathbf{x}})}\ .$$

The Mahalanobis distance weights the distance of a multi-dimensional data point $\mathbf{x}$ from its mean $\bar{\mathbf{x}}$ such that observations that are on the same multivariate normal density contour will have the same distance d_m. These Mahalanobis iso-distance contours take the form of ellipses in two-dimensional space, and ellipsoids in three-dimensional space. Obviously, the number of points in the profile $2k+1$ is normally much higher than three, in which case the iso-distance loci take the form of hyper-ellipsoids in multi-dimensional space.

The gray level appearance model is obtained from a set of training images, and consists for each landmark in a mean profile and a covariance matrix for each of the features. The total of stored models is thus proportional to the number of landmarks and the number of features.

Shape models

The positional correlation between landmarks along the contour is taken into account in a shape modelling step. Because segmentation of anatomic structures with relatively deterministic shape is aimed at, knowledge regarding the shape must preferably encode the deterministic component, along with relevant anatomical variation, and further preferably ignore irrelevant shape details.

Two embodiments are outlined to take into account positional correlation. The first is based on a global approach involving the use of all landmark positions simultaneously; the second employs a local approach, modelling the spatial disposition of successive landmarks.

PCA analysis and fitting

In the first embodiment, a variance-covariance matrix of landmark coordinate positions is constructed from manually placed landmarks in a set of aligned training images. Each diagonal entry in the variance-covariance matrix represents the variance of the coordinate values of the landmarks. The non-diagonal entries represent the co-variance of pairs of landmark coordinate values, i.e. it expresses how the coordinates of two landmarks vary with respect to each other. The variation may systematically be in equal 'positive directions' with respect to their means, yielding high positive covariance, or systematically in equal 'negative directions' with respect to their respective means, yielding high negative covariance, or randomly in opposite directions with respect to their means, yielding zero covariance. An eigenvector/eigenvalue analysis, known in the prior art as Principal Components Analysis (PCA), will detect any global systematic variation modes of the landmarks with

respect to the mean landmark positions. PCA aims to reduce the dimensionality of a data set by only keeping the components of coordinate set with largest variation; it therefore statistically decorrelates the original coordinate positions. By projecting onto these highly varying subspaces, the relevant statistics can be approximated by a smaller dimensional system that retains important shape features. A sufficient number of these principal modes are retained and used as the shape model, ignoring the other modes attributed to noisy variations.

Landmark connection vector model

The second embodiment aims at modeling the disposition of successive landmarks at the local level. However, this restriction does not imply that the global shape is not captured. The shape outline can be completely reconstructed given the start point coordinates and the connection vectors, by iteratively adding the connection vector to the current point, starting from the start point. This property is exploited in shape description at the pixel level by chain codes for example. The connection vector is a discrete approximation of the local tangent vector to the curve. When taking into account the direction vector between successive landmarks, i.e. first order geometric derivatives, it allows describing the shape of the object irrespective of translations of the object. By taking into account the curvature, i.e. second order geometric differentiation, between successive landmarks, for example expressed by the angle subtended by two successive connection vectors, the shape of the object can be described irrespective of translations and rotations of the shape (so-called rigid body transformations). In that case the curve can be completely reconstructed when the start position, and the first geometric derivative (i.e. start connection vector) is known, by iteratively adding the difference vector of successive tangent vectors. Curvature based on a chain code at the pixel level is usually highly noise dependent for the scale is at a too fine level. This drawback is coped with in the present disclosure because the landmarks are sufficiently spaced apart so as to enable computation of noise-independent yet accurate first or second order derivative

information. Because the magnitude of tangent or curvature vector depends on scale, normalizing the size of each curve of the training set to unity for example, further allows describing the curve up to an Euclidean similarity transformation. The mean and covariance statistics of displacement vectors (i.e. first derivative or tangent) or vector differences of displacement vectors (i.e. second derivative or curvature) between manually placed landmarks in a set of aligned training images constitutes the shape model of this embodiment.

Segmentation systems

The above-described models can be used in a segmentation system.

Segmentation systems can be organised along two tracks each consisting of building blocks with related objectives.

A first building block aims at confining the admissible positions of instantiated model landmark positions to a number of most probable locations in a target image.
In the case landmarks are represented by linear profiles in feature images, a best point voting approach can be adopted that ranks a set of candidate positions on a perpendicular line through an instantiated landmark.
For a circular profile case, a rectangular search grid can be considered at each instantiated landmark and grid crossing positions are ranked according to the combined Mahalanobis distance of circular profiles through all feature images to the model circular profiles.
Second and third building blocks jointly constitute an optimization step. A cost matrix may be constructed, using the cost measures associated with the retained admissible position for each landmark. The cost comprises a term associated with agreement in gray level appearance. Agreement with the shape model may be verified at the global level using weighted PCA contour fitting. Alternatively, the shape agreement may be verified by incorporating a shape cost

associated with the candidate position of the landmark. A dynamic programming algorithm can be used to construct a minimum cost path through the cost matrix, yielding the final segmentation as the path with lowest overall cost.

Finally, characteristic points on the final segmentation may indirectly define the geometrical position of the key measurement points, and all depending measurement objects and measurement quantities may be generated on the basis of these key measurement points.
Alternatively, the landmark points used in the segmentation may be used to directly define the key measurement points and a search grid for each measurement point may be provided to locate the admissible positions of the measurement point.

Further specific embodiments and associated advantages of the present invention will become apparent from the following description and drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a flow chart illustrating a first modeling system (Modeling system I),
Fig. 2 is a flow chart illustrating a first segmentation system (Segmentation System I),
Fig. 3 is a flow chart illustrating a second modeling system (Modeling System II),
Fig. 4 is a flow chart illustrating a second segmentation system (Segmentation System II),
Fig. 5 is a flow chart illustrating measurement computation on the basis of the characteristic points on the anatomic objects segmented making use of a segmentation method in which models generated according to the present invention are used,
Fig. 6 is a measurement scheme for measurement of the cardiothoracic ratio from a chest radiograph,

Fig. 7(a) illustrates manual delineation of the left lung in a chest radiograph, in fig. 7(b) the left lung is represented as a set of landmarks $\{p_i\}_{i=1}^{n}$, in this example $n=14$.

Fig. 8 illustrates a discrete representation of the lungs fields by a set of n landmarks $\mathbf{p}_1=(x_1,y_1),...,\mathbf{p}_n=(x_n,y_n)$. The image sizes are normalized between 0…1. The continuous curve is interpolated through the discrete landmarks.

Fig. 9 illustrates a linear and a circular gray-level profile as a sampling of gray values at a number of points on a line resp. a circle with radius r_c around the landmark,

Fig. 10 shows experimentally estimated distributions for some connection vectors on the left lung shape consisting of 14 landmarks: (a) $\mathbf{v}_1$, (b) $\mathbf{v}_2$, (c) $\mathbf{v}_3$.

Fig. 11 shows first moment feature images of a chest radiograph, corresponding to the derivatives up to second-order $\left(L,L_x,L_y,L_{xx},L_{yy},L_{xy}\right)$ and the scales σ equal to 0.5, 2 and 8 pixels,

Fig. 12 shows second moment feature images of a chest radiograph, corresponding to the derivatives up to second-order $\left(L,L_x,L_y,L_{xx},L_{yy},L_{xy}\right)$ and the scales σ equal to 0.5, 2 and 8 pixels,

Fig. 13 shows examples of application of the present invention to the segmentation of lung fields on a chest radiograph: (a) manual segmentation of left and right lung field; (b) automatic segmentation of left and right lung field.

Fig. 14 illustrates the computation of the Cardiothoracic Ratio using (a) the manual lung field segmentations resulting in a $\mathrm{CTR}_{\mathrm{man}}=0.47$; (b) using the automatically derived lung field segmentations resulting in a $\mathrm{CTR}_{\mathrm{automatic}}=0.45$.

DETAILED DESCRIPTION OF THE INVENTION

The present invention will be explained in detail with regard to a specific application, namely the segmentation of the lung field in a medical image.

Object representation

In the specific embodiments of the method of the present invention that are described below, an anatomical object in an image is represented mathematically as a fixed number of discrete, labeled points lying on the contour that encloses the object, i.e. $p_1 = (x_1, y_1), ..., p_n = (x_n, y_n)$.

The contour $\{p_i\}_{i=1}^{n}$ runs from p_1 to p_n and back to p_1. Hence, the object may be captured by a discrete shape vector $x = (x_1, y_1, ..., x_n, y_n)^T$. The coordinate system is chosen such that all points inside the image area lie in the domain [0,1]x[0,1] (fig. 7).

Additionally, characteristic points lying in the region enclosed by the contour may be added to the object representation as well, allowing e.g. measurement of entities lying inside the object.

The segmentation scheme described below needs a number of training images, in which the shape vectors x are manually determined. Once the algorithm is trained on the data set, it is able to produce the shape vector in a new image.

The continuous contour of an anatomical object can be reconstructed by curve interpolation from the points $\{p_i\}_{i=1}^{n}$. A first order interpolation links successive points with a line segment, resulting in a polygonal approximation. Higher order interpolations may be used to achieve a smoother outline. Irrespective of the order, the continuous representation can be used to derive spatial properties of the curve, such as the tangent direction along the curve, or the normal direction at a certain point on the curve.

In three dimensions, the third axis z is added, and the landmark points are now distributed on a surface that encloses the object, i.e. $p_1 = (x, y_1, z_1), \ldots, p_n = (x_n, y_n, z_n)$. Hence, in this point-based representation, the object may be captured by a discrete shape vector $x = (x_1, y_1, z_1, \ldots, x_n, y_n, z_n)^T$.

A first order continuous representation is achieved by linking groups of surface points in polygons, each polygon defining one interpolating surface patch of a polygon mesh, connected such that each edge is shared by at most two polygons. An edge connects two vertices, and a polygon is a closed sequence of edges. When the points are grouped in triples, each triple defines exactly the equation of a planar surface patch, in which case the polygonalization process becomes a surface triangulation. This process is different from methods in the prior art that extract iso-surfaces and locally triangulate the surface (examples of these methods are the Marching Cubes algorithm introduced in W. Lorensen, H.E. Cline, Marching Cubes: a high resolution 3D surface construction algorithm, *Computer Graphics*, Vol. 21, No. 4, July 1987, p. 163-169, and later refinements known as the NOISE, WISE and SAGE algorithms). The difference lies in the fact that the 3D object representation adopted here does not assume that the points have equal gray value, hence the object shape or organ outline is assumed to be different from an iso-gray value surface hence the segmentation cannot be based on iso-surface extraction. Instead, the ensemble of landmarks will be modeled geometrically and radiometrically and the models are applied in a novel model-based 3D segmentation scheme.

Additionally, characteristic landmark points lying in the volume enclosed by the surface may be added to the object representation as well, thereby (a) enabling influencing and steering of the segmentation result by dominant anatomic features inside an organ or (b) allowing e.g. measurement of entities lying inside the object. A polyhedral mesh may now be constructed linking

points on the surface and inside the surface into an arrangement of polyhedral volumetric cells. In 3D lung scans for example, these additional points may comprise characteristic points of the vessel trees and bronchial trees (such as bifurcation points). In particular the polyhedron can comprise only four points, each group forming a tetrahedron, and the polyhedralization now becomes a tetrahedralization. The outer shell of the tetrahedral structure representing the volume is a triangulated surface.

Modeling System I (fig. 1)

Gray level appearance model

The gray level appearance model (also called photometric model) captures the spatial intensity distribution in the neighborhood of so-called landmarks (also called landmark points). The spatial component of the gray level appearance model has been devised along the use of gray-level profiles. As will be apparent from their construction, other schemes can be employed.

The number n of landmarks that is for example used to accurately describe a lung field outline may be as low as 14, e.g. landmark 1 denotes the uppermost point of each lung field, landmark 7 is positioned at the cross section of either the left or right heart shadow with their respective hemi-diaphragm, landmark 9 denoting the costophrenic angle. The other landmarks can be equidistantly positioned between these major landmarks. Obviously, a higher number of landmarks may be used.

- Linear gray-level profiles (fig. 9)

At each landmark, the gray level appearance model describing the typical image structure around the landmark is sampled perpendicular to the contour. On either side of the landmark, k pixels are sampled using a certain step size, which gives profile vectors of length $2k+1$. This sampling scheme has the advantage that it is able to

model linear intensity gradients at the landmark. The profiles may be normalized so that the sum of absolute values of elements in the profile is 1. The direction of sampling can be computed as the mean normal direction between the normals on the line segments resulting from connecting the landmark with its previous or its next neighbour, or as the normal on the line segment connecting the previous and next landmark.

In three dimensions the computation of the normal direction is more compound for the normal to a surface represented as a distributed point set needs be found. In three dimensions, the normal is defined with respect to a plane. This plane may be obtained by fitting a plane $ax+by+cz+d=0$ through a set of at least three selected nearest neighboring points, using Euclidean distance as the closeness criterion. The normal to the resulting plane has components (a_i,b_i,c_i), and is used as the direction vector along which the volume is sampled perpendicular to the object surface at a particular landmark (x_i,y_i,z_i) to obtain a 3D voxel-level profile.

When landmark points inside the volume are added to the model, profiles may be sampled in the volume for example along some of the connection vectors of the landmark point with landmarks points surrounding it and that are lying either inside the surface or on the surface.

- Gray-level feature images

The image structure around a landmark exhibits is captured in a number N of mathematical features as explained higher. The feature images are obtained by computing a number of moments of the local histogram of derivative images in a window with width α around each location x_0. In a preferred embodiment the first and second histogram moments, all derivatives up to second-order $\left(L,L_x,L_y,L_{xx},L_{yy},L_{xy}\right)$ and 5 inner scales (σ=0.5,1,2,4,8 pixels) are

computed, amounting to a total of $N = 2 \times 6 \times 5 = 60$ feature images. The window scale to compute the histogram is in accordance to the inner scale, i.e. $\alpha = 2\sigma$. The set of feature images can be regarded as a gray-level decomposition of the gray-level function in the neighbourhood of a landmark.

An example of the computed feature images in a thorax image is given in fig. 11 (first histogram moments of original image and 5 derivatives images at three scales) and fig. 12 (second histogram moments of original image and 5 derivatives images at three scales). Thorax images normally appear in the upright position in the image; anatomic objects pertaining to other radiographic examinations may exhibit non-standardized position and rotation, in which case the construction of feature images may employ Cartesian Differential Invariants.

In three dimensions, the derivatives are 3D operators. The volumetric data can locally be described by images containing the partial derivatives of the original image $\left(L, L_x, L_y, L_z, L_{xx}, L_{yy}, L_{zz}, L_{xy}, L_{yz}, L_{xz}, \ldots\right)$. Similar to the 2D case, the derivatives may be computed at a number of blurring scales σ , and finally be used to compute a number of feature images using moments of the local histogram. In a preferred embodiment, the first and second histogram moments, all derivatives up to second-order $\left(L, L_x, L_y, L_z, L_{xx}, L_{yy}, L_{zz}, L_{xy}, L_{yz}, L_{xz}\right)$ and 5 inner scales (σ =0.5,1,2,4,8 pixels) are computed, amounting to a total of $N = 2 \times 10 \times 5 = 100$ feature images. The window scale to compute the histogram is again in accordance to the inner scale, i.e. $\alpha = 2\sigma$.

- Gray-level appearance model

All images of a training set are used to build a statistical model for each feature and each landmark. Denoting the normalized profile sampled at a landmark i (e.g. 14 in total for the 2D case of a

thorax lung field) in feature j (e.g. 60 in total for 2D images) as a vector $\mathbf{g}_{i,j}$, the probability distribution of $\mathbf{g}_{i,j}$ is estimated from the training images, resulting in a mean $\bar{\mathbf{g}}_{i,j}$, and a covariance $S_{i,j}$. For a length of a profile vector of $2k+1$, the size of this covariance matrix is $(2k+1)\times(2k+1)$. The gray level appearance model for a certain landmark i consists of the mean profiles $\bar{\mathbf{g}}_{i,j}, j=1...N$ and covariance matrices $S_{i,j}, j=1...N$. The gray-level appearance model for the total lung field consist of the collection of all mean profiles $\bar{\mathbf{g}}_{i,j}, i=1...n, j=1...N$ and covariance matrices $S_{i,j}, i=1...n, j=1...N$.

In three dimensions, voxel-level samplings (also termed 'profiles') for each of the n 3D landmarks are obtained in each of the N feature images. The 3D statistical model of the landmarks consists of the collection of all mean profiles $\bar{\mathbf{g}}_{i,j}, i=1...n, j=1...N$ and covariance matrices $S_{i,j}, i=1...n, j=1...N$.

- Gray-level cost

The Mahalanobis distance between a new profile $\mathbf{g}_{i,j}$ sampled at a point $\mathbf{p}$ in a new feature image j for a certain landmark i is computed as

$$h_{i,j}(\mathbf{p}) = \left(\mathbf{g}_{i,j}(\mathbf{p}) - \bar{\mathbf{g}}_{i,j}\right)^T S_{i,j}^{-1} \left(\mathbf{g}_{i,j}(\mathbf{p}) - \bar{\mathbf{g}}_{i,j}\right)$$

A smaller Mahalanobis distance means a larger probability that the profile $\mathbf{g}_{i,j}(\mathbf{p})$ originates from the Gaussian distribution with mean $\bar{\mathbf{g}}_{i,j}$ and covariance $S_{i,j}$. Hence, the Mahalanobis distance may be used as a gray-level cost measure, denoted as $h_{i,j}(\mathbf{p})$. This cost

measure is a function of the location $\mathbf{p}$ in a given image. The location that most likely is the true location for the landmark i, in accordance to feature j, is the point $\mathbf{p}$ that minimizes $h_{i,j}(\mathbf{p})$. This gray-level cost can thus be defined for each feature j.

In three dimensions, the cost function at a certain 3D point $\mathbf{p}$ is likewise evaluated as the Mahalanobis distance $h_{i,j}(\mathbf{p})$ between a new 3D voxel-level profile $\mathbf{g}_{i,j}$, sampled at the 3D point $\mathbf{p}$ in a new 3D feature image j for a certain 3D landmark i, and the mean $\bar{\mathbf{g}}_{i,j}$.

Shape model

The curve outline, in the case of lung fields: one for each of the lung fields, is described by n landmark points (fig. 8). These landmark points are manually determined in a set of s training images, yielding a sequence of points $(x_1, y_1)...(x_n, y_n)$. These coordinate tuples are subsequently positioned in a vector $\mathbf{x} = (x_1, y_1, ..., x_n, y_n)^T$, representing the curve. Next Principal Components Analysis (PCA) is applied to the shape vectors $\mathbf{x}$ of the training images. The PCA projects the curve in a lower dimensional space, covering the most important modes of geometric variation of the lung fields. Each curve $\mathbf{x} \in \mathbb{R}^{2n}$ can be approximated by $\mathbf{b} \in \mathbb{R}^{t}$ with $t \leq 2n$:

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \mathbf{b}$$

with $\bar{\mathbf{x}}$ the mean shape and $\Phi \in \mathbb{R}^{2n \times t}$ the eigenvectors of the covariance matrix of $\mathbf{x}$ corresponding to the t largest eigenvalues. Each eigenvalue determines the variance of the shape for that shape mode. This curve approximation, represented by the vector $\mathbf{b}$, constitutes the shape model and fits the curve $\mathbf{x}$ into the shape model. The eigenshapes can be regarded as a zero-order (positional) geometric decomposition of the shape represented by the set of landmarks. The role of the shape model is to constrain the deformation of the

object between the bounds of variation imposed by the training set, for example the three standard deviations with respect to the mean shape.

In three dimensions, the surface outline (i.c. one outline for each anatomic object) is described by n 3D points (landmarks). The outlines may for example represent the 3D lung fields. These landmarks may be manually determined in a set of s training images, or obtained with bootstrapping techniques from a previous sub-set $s'$of the training images, yielding a sequence of 3D points $(x_1, y_1, z_1)...(x_n, y_n, z_n)$. These coordinate triples are subsequently positioned in a vector $\mathbf{x} = (x_1, y_1, z_1, ..., x_n, y_n, z_n)^T$, representing the surface. As noted earlier, additional characteristic points inside the volume delineated by the surface may complement the shape model.

Next, Principal Components Analysis is applied to set of 3D landmarks covering the surface of the object(s) in each 3D training image. The PCA analysis projects the surface in a lower-dimensional space, covering the most important modes of geometric variation of the anatomical object(s) (e.g. the 3D lung fields). Each surface $\mathbf{x} \in \mathbb{R}^{3n}$ can be approximated by $\mathbf{b} \in \mathbb{R}^{t}$ with $t \ll 3n$:

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \mathbf{b}$$

with $\mathbf{x}$ the mean shape and $\Phi \in \mathbb{R}^{3n \times t}$ the eigenvectors of the covariance matrix of $\mathbf{x}$ corresponding to the t largest eigenvalues. Each eigenvalue determines the variance of the shape for that shape mode. This surface approximation, represented by the vector $\mathbf{b}$, constitutes the shape model and fits the surface $\mathbf{x}$ into the shape model. Similar to the 2D case, the eigenshapes can be regarded as a zero-order (positional) geometric decomposition of the surface shape represented by the set of landmarks. The role of the surface shape model is to constrain the deformation of the object between the

bounds of variation imposed by the training set, for example the three standard deviations w.r.t. the mean shape.

Segmentation System I (fig. 2)

The segmentation algorithm will initially position a curve, representing a lung field outline, in the image, and finds a segmentation solution by iteratively updating the positions of the curve's constituent landmarks until an optimization criterion reaches an optimal value. Each iteration consists of the following steps.

Step 1. Best point voting

Assuming a current position of a particular landmark, a better location for this landmark will be searched on a line perpendicular to the contour. On each side of the current position of the landmark, n_s positions are evaluated. One of the $2n_s+1$ points will be voted as the best next point. Each feature of the set of N features contributes one vote to one of the $2n_s+1$ points. That point of the $2n_s+1$ points is selected that has minimum cost according to the Mahalanobis distance criterion. For each of the $2n_s+1$ points, the profile consisting of $2k+1$ points sampled in feature image j is put in the equation of $h_{i,j}(\mathbf{p})$ with the proper mean $\bar{\mathbf{g}}_{i,j}$ and covariance $S_{i,j}$ filled in. The point with the lowest Mahalanobis distance is voted by that feature. Each feature chooses one point, resulting in N votes divided over $2n_s+1$ points as sets $N_1,\ldots,N_{2n_s+1}$ with $\sum_{k=1}^{2n_s+1} N_k = N$. Obviously the number N of features must be substantially larger than the number $2n_s+1$ of selectable points for otherwise there are too few votes that may be assigned to a point of the set of selectable points.

The number of votes $N_{i,j}$ that each point of the set of considered $2n_s+1$ points for a certain landmark i receives, may be

regarded as a confidence rate that the particular point is the searched for landmark.

In three dimensions, the choice among the $2n_s+1$ 3D points is obtained by taking the maximally voted point through the three-dimensional features.

Step 2. Minimal cost path optimization

According to the first step, each landmark may separately be displaced over $2n_s+1$ possible locations. The confidence rates $N_{i,j}$ can be put in a matrix as follows

$$R=\begin{bmatrix} N_{1,1} & \cdots & N_{1,2n_s+1} \\ \vdots & N_{i,j} & \vdots \\ N_{n,1} & \cdots & N_{n,2n_s+1} \end{bmatrix} \quad \text{with} \quad \sum_{j=1}^{2n_s+1} N_{i,j}=N \quad \forall i=1,\ldots,n \quad .$$

Updating the contour involves choosing one element in each row of the matrix C; the concatenation of all elements forming a path through the matrix, such that a cost criterion is optimized (i.c. minimized), is a minimal cost path (MCP) search.

A logical choice would be to take the element with highest confidence in each row; this approach has the drawback that it does not take into account the presence of outliers. Therefore, it may include abrupt changes in the path. This drawback is overcome by imposing a smoothness constraint that excludes the abrupt changes.

To this purpose, first a cost matrix is constructed as follows:

$$C=\begin{bmatrix} 1/N^m_{1,1} & \cdots & 1/N^m_{1,2n_s+1} \\ \vdots & 1/N^m_{i,j} & \vdots \\ 1/N^m_{n,1} & \cdots & 1/N^m_{n,2n_s+1} \end{bmatrix} \quad ,$$

with the power m in the denominator being a positive number that influences the computed path (a larger m yielding smaller cost contribution for a point and hence steering the path towards this point). Second, a cost function is constructed as follows:

$$J(j_1, j_2, ..., j_n) = \sum_{i=1}^{n} C_{i,j_i} + \alpha \sum_{i=1}^{n} \left| \delta_{i,j_i} - \delta_{i-1,j_{i-1}} \right|,$$

with δ_{i,j_i} the displacement of point i towards its target point j_i along the normal to the shape and α a weight factor. The second term in this equation is the smoothness constraint that penalizes target point configurations for which displacements of one target point w.r.t. the displacement of the previous target point vary abruptly along the contour. The cost function J is optimized, using e.g. a technique known in the prior art as dynamic programming, described in e.g. D.H. Ballard, C.M. Brown, Computer Vision, Englewood Cliffs, Prentice Hall Inc. 1982, pp 137-143, resulting in the target points $(j_1^*, ..., j_n^*)$, that minimize J as a minimal cost path (MCP).

Alternatively to the matrix data structure, a layered graph may be used also. Each layer represents the array of perpendicularly sampled image points at a landmark. The graph is augmented with a set of arcs from the last layer back to the first. The dynamic programming technique extracts the boundary as an optimal closed path through the graph.

In three dimensions, the problem of solving for the minimum cost path is more compound and not straightforwardly derivable from the 2D case as will be apparent from the sequel. First, the matrix data structure that is constructed in the 2D contour detection approach cannot be extended for the 3D case because there is no obvious way to re-circulate (in the 2D case this is implemented by copying the first row after the last row in the matrix). Furthermore the 2D matrix method enables a systematic way to visit all landmarks

(each corresponding to a row); a straightforward way is a sequential order by concatenating successive landmarks. In 3D, each landmark has a number of nearest neighbors spanned over the 3D surface that, in order to treat those on an equal basis must be visited in parallel. Second, the number of variables to minimize is much larger because the surface of an object in a 3D image comprises usually much more points than a contour of an object in a 2D image. Therefore, the minimization algorithm may get more easily stuck in a local minimum for not all combinations of positions may be exhaustively evaluated. The simulated annealing optimization strategy is outlined to cope with this problem.

Two optimization strategies suitable for 3D surface optimization are outlined in the sequel.

Minimal cost surface optimization

1. A weighted undirected graph data structure G is created connecting all landmark points on the surface departing from the triangulated surface, and possibly augmented with connections towards points inside the surface, forming an overall polyhedral mesh. Each node (vertex) g comprises the 1D array of points and the associated cost (e.g. inverse gradients) of the points. The points in the points array are sampled perpendicular to the surface towards the inside and outside of the surface. Each node is linked to all its neighboring nodes obtained by following the edges (arcs) a departing from the node. Any node that has already been visited previously is not visited anymore to avoid cycles. The edges of the graph may be weighted with (a) the calculated Euclidean distance between the point pair that is linked by the edge (b) a geometric property of a landmark and its neighboring landmarks such as the local curvature. The resulting graph is sparse for there exist no edges between all non-neighboring landmark points. Alternative representations may be used such as an adjacency matrix representation where the weights of an existing adjacency relation between a pair of points are put as

an entry in the matrix; all non-neighboring points receive a non-existent weight. Another alternative representation for the graph is the adjacency list, which is more suited when the graph is sparse, which is the case for graphs created from a surface triangulation. Here, a linked list is maintained for each landmark storing all neighboring points in subsequent cells of the list; cross links between these lists may be provided to algorithmically ease e.g. the deletion of a node.

The collection of nodes each containing a 1D point array may be viewed as a shell of points around the current surface. A new surface has to be computed comprised of the calculated new locations in the array for each node. This process is outlined in the sequel steps and may be viewed as a deformable surface segmentation.

2. Starting from a node g in the graph, visit all other nodes of the graph in a systematic way. Graph algorithms, described in standard text books on algorithms such as R. Sedgewick, Algorithms in C++, Addison-Wesley, 1992, are applied to this purpose.

a. Minimum spanning tree (MST)

This step is solved using methods of the prior art known as minimum spanning trees. A MST of a weighted graph is a collection of edges connecting all the nodes such the sum of interconnecting weights of the edges is at least as small as the sum of the weights of any other collection of edges connecting all the nodes.

The solution of this problem is based on a property of graph theory in that given any division of the nodes of a graph into two sets, the MST contains the shortest of the edges connecting a node in one of the sets to a node in the other set.

The MST is built starting from any node (vertex), and always takes the next node that is “closest” (resulting in an Euclidean MST, but extremality may be according to other criteria) to the

nodes already taken. In other words, the edge of lowest weight is followed among all those edges that connect nodes already in the tree to nodes not yet in the tree, and the node that the edge leads to is added to the tree. The aforementioned property of minimum spanning trees guarantees that each edge added is part of the minimum spanning tree. The resulting MST need not be unique. Practical methods to find the MST are known in the prior art as the method of Kruskal or Prim's algorithm.

b. Breadth first traversal

Another way of visiting the nodes of the graph in a systematic way is known in the prior art as the breadth first traversal starting from a certain node. Breadth first search starts covering the area of the surface around a given point node, moving farther away on the surface only when everything close has been looked at. This kind of graph traversal may be regarded at propagating away from the starting point so as to encompass gradually all parts of the surface. For closed surfaces, i.e. surfaces without boundaries, parts of the propagating front will collide with each other at the backside of the surface.
When all weights are one in the graph G, the breadth first traversal tree leads to a shortest path spanning tree.
In general, the MST tree will be different from the breadth first tree.
Although the distances change between each node and its neighboring node and hence the MST or breadth first tree must be recomputed, these trees may be kept constant during all iterations of the dynamic programming algorithm of the next step.
It should be noted that although the initial graph will in general be not planar, the resulting visiting MST and breadth first trees are.
In this way all nodes of the segmented surface and its complementary inner landmarks are linked together in a systematic way.

3. The dynamic programming (DP) approach is now performed on the MST or breadth first traversal tree, using the gradient in each sampled point and the difference in displacement of a node and the displacements of the neighboring nodes, e.i.

$$J(j_1, j_2, ..., j_n) = \sum_{i=1}^{n} C_{i,j_i} + \alpha \sum_{i=1}^{n} \left| \delta_{i,j_i} - \delta_{i-1,j_{i-1}} \right|$$

with n the number of nodes (points) in the tree, δ_{i,j_i} the displacement of point i towards its target point j_i along the normal to the surface, $\delta_{i-1,j_{i-1}}$ the displacement of the parent node point in the tree, and α a weight factor. Node i and $i-1$ represent neighboring points on the surface; j_i and j_{i-1} are the indices in the point array associated with nodes i resp. $i-1$. The second term in this equation is the smoothness constraint that penalizes target point configurations for which displacements of one target point w.r.t. the displacement of the previous target point vary abruptly over the surface. The cost function J is optimized, using e.g. a technique known in the prior art as dynamic programming, resulting in the target points $(j_1^*, ..., j_n^*)$, that minimize J as a minimal cost path (MCP). When inner points are added to the model, displacement for these nodes may be along one of the connecting vectors with surrounding nodes.

Starting from the positions obtained for the landmarks, steps 2 and 3 may now be iterated starting from another node g in the graph, from which a new MST is grown and new positions of landmarks are computed until a stopping criterion is met, e.g. the sum of absolute displacements of all nodes between two iterations does not change noticeably anymore.

Simulated annealing optimization

Simulated annealing has been introduced in the prior art by S. Kirckpatrick, C.D. Gelatt, M.P. Vecchi, Optimization by simulated annealing, Science, Vol. 220, No. 4598, May 1983, p. 671-680. In the sequel, details pertaining to the implementation according to the

present invention are outlined. In this optimization strategy, no tree need be constructed from the graph to visit the points on the surface in a systematic way. With each ensemble of positions of landmarks on the current surface, the energy associated with it can be calculated from the graph G as follows

$$E = \sum_{i=1}^{n} \left(C_{i,j_i} + \alpha \sum_{n_i \in N_i} \left| \delta_{i,j_i} - \delta_{n_i,j_{n_i}} \right| \right).$$

The energy is the objective function that has to be minimized. It is composed of a term associated with the node i and a term associated with the set N_i of neighborhood nodes of the node i. The node term assigns a lower energy when the point has higher local gradient for example. The neighborhood term assigns lower energy e.g. when the displacement of each neighboring node w.r.t. the displacement applied to the current node is smaller, meaning a surface shape is penalized when it deforms so as to create local abrupt changes. The simulated annealing optimization consists of applying the Metropolis algorithm in the calculation of the optimal surface: the energy always takes the option for a lower energy state (a downhill step) when the energy E' of the new iteration is lower than the energy E of the previous iteration, while sometimes taking an uphill step when the energy E' is higher than energy E. The probability of such change from configuration of energy is adopted from thermodynamic systems and obeys the so-called Bolzmann probability distribution $p = \exp\left[-(E'-E)/kT\right]$. The control parameter T (analog of temperature) and an annealing schedule tells how the energy E is lowered from high to low values, e.g. after how many random changes in configuration is each downward step in T taken, and how large is that step. In contrast to the tree-based methods outlined before, no systematic way of visiting the points on the surface is taken; instead they are taken at random over the surface. The steps of the SA algorithm are as follows

1. Set T to a large value.

2. (Configuration step) Initialize the surface as a point set $(x_1, y_1, z_1)...(x_n, y_n, z_n)$ and compute its energy E.

3. (Rearrangement step) Deform the surface by generating a new position $(x_1', y_1', z_1')...(x_n', y_n', z_n')$ for the landmarks along their direction perpendicular to the surface. The search algorithm may either opt to displace one randomly chosen landmark at a time applying a certain step size, or displace all of them simultaneously. Efficient set of moves may be obtained by tracking the sense (surface inward or surface outward) in which the move of each point occurs and continuing when it is in the direction of lower energy.

4. (Objective function step) The energy E' of the resulting surface is computed.

   If $E' < E$ accept the change, go to 6.

5. (Annealing schedule) Compute $p = \exp[-(E' - E)/kT]$.

   If $p > P$, then accept the change, i.e. although the new surface involves an energy increase w.r.t. the previous surface, the probability is sufficiently high to accept the uphill step. P is a random number between 0 and 1.

6. (Annealing schedule) Decrease T slightly. If $T > T_{\min}$ go to 3, else stop.

As this step gradually lowers the temperature, the uphill moves become gradually less likely until for low values of T such moves become impossible, meaning that the solution starts to settle in towards the final segmentation, only allowing small refinements.

Because the tree-based methods follow a systematic path through the points on the surface, and because no unique tree exists, the solution for the surface segmentation may be suboptimal. The simulated annealing method has no pre-defined path, enabling a greater chance of getting into the global optimum, however at the expense of requiring more computation time.

Step 3. Unweighted and weighted PCA contour fitting

In the previous step, each landmark i is given a new location j_i^*. The sequence $(j_1^*,...,j_n^*)$ represents a contour that may be fitted in the shape model $\mathbf{x} \approx \bar{\mathbf{x}} + \Phi\mathbf{b}$ by substituting it for $\mathbf{x}$ and deriving $\mathbf{b}$ by solving $\Phi\mathbf{b} = \mathbf{x} - \bar{\mathbf{x}}$. Fitting $\mathbf{x}$ in the shape model means finding a $\mathbf{b}$ value such that $\bar{\mathbf{x}} + \Phi\mathbf{b}$ approximates $\mathbf{x}$ with small errors in the high confidence landmarks.

Un-weighted PCA contour fitting

Since $\Phi \in \mathbb{R}^{2m \times t}$ there are more equations than unknowns, and the system is over-determined, having no exact solution. Therefore, $\left\| \Phi\mathbf{b} - (\mathbf{x} - \bar{\mathbf{x}}) \right\|_p$ must be minimized for some suitable choice of p. Different norms render different optimum solutions. In particular the case of $p = 2$, the least squares problem, is more tractable, and is solved in the prior art by methods based on normal equations and the QR factorization. Denoting the difference between actual point sequence and the fitted point sequence as the fit error $\mathbf{e} = \mathbf{x} - (\bar{\mathbf{x}} + \Phi\mathbf{b})$, the problem is to seek the $\mathbf{b}$ to minimize $\mathbf{e}^T\mathbf{e}$. Substituting for $\mathbf{e}$, this optimization problem becomes

$$\min_{\mathbf{b}} J(\mathbf{b}) = \min_{\mathbf{b}} \mathbf{e}^T\mathbf{e} = \min_{\mathbf{b}} ((\mathbf{x} - \bar{\mathbf{x}})^T - \mathbf{b}^T\Phi^T)((\mathbf{x} - \bar{\mathbf{x}}) - \mathbf{b}\Phi) .$$

$J(\mathbf{b})$ is minimized if $\nabla J(\mathbf{b}) = 0$. The gradient $\nabla J(\mathbf{b})$ of $J(\mathbf{b})$ can be written

$$\nabla J(\mathbf{b}) = -2\Phi^T W^2(\mathbf{x} - \bar{\mathbf{x}}) + 2\Phi^T W^2\Phi\mathbf{b} .$$

Hence the optimization problem leads to the equations

$$\nabla J(\mathbf{b}) = 0 \Leftrightarrow \Phi^T\Phi\mathbf{b} = \Phi^T(\mathbf{x} - \bar{\mathbf{x}}) ,$$

known as the normal equations. These are solved for $\mathbf{b}$ as

$$\mathbf{b} = (\Phi^T \Phi)^{-1} \Phi^T (\mathbf{x} - \bar{\mathbf{x}})$$

This equation projects the data vector $\mathbf{x}$ of the image space into the PCA model space and returns the approximation coefficients $\mathbf{b}$. The approximation result $\mathbf{b}$ may be projected back in image space by evaluating the shape model $\mathbf{x} = \bar{\mathbf{x}} + \Phi \mathbf{b}$.

Weighted PCA contour fitting

The corresponding number of votes N_{i,j_i^*} may be used as weights when fitting the contour into the shape model. As a consequence, the PCA model will tend to give a high priority to the high confidence points. Points with a low confidence rate will lessen their influence on the contour that the model generates. The optimization problem now consists of minimizing the weighted quadratic error $(W\mathbf{e})^T (W\mathbf{e})$. Using a similar approach with normal equations as in the un-weighted case, this problem leads to the solution for $\mathbf{b}$ expressed as

$$\mathbf{b} = (\Phi^T W^2 \Phi)^{-1} \Phi^T W^2 (\mathbf{x} - \bar{\mathbf{x}}) .$$

This equation is evaluated to project $\mathbf{x}$ in the PCA space and may be projected back using the shape model $\mathbf{x} = \bar{\mathbf{x}} + \Phi \mathbf{b}$.

Steps 1 to 3 are iterated until a stop criterion is fulfilled, e.g. when the positions of the landmarks do not change substantially between successive iterations.

Both the un-weighted and weighted PCA fitting yields the final segmentation result as a set of landmark points $\mathbf{x}$ in the image. A continuous spline curve may be interpolated through these points to represent the segmentation analytically.

Modeling System II (fig. 3)

Gray level appearance model

The gray level appearance model captures the spatial intensity distribution in the neighborhood of landmarks. In this embodiment, the spatial model departs from circular profiles sampled in the neighborhood of a landmark. The number n of landmarks that is used to accurately describe a lung field outline may be as low as 14, e.g. landmark 1 denotes the uppermost point of each lung field, landmark 7 is positioned at the cross section of either the left or right heart shadow with their respective hemi-diaphragm, landmark 9 denoting the costophrenic angle. The other landmarks are equidistantly positioned between these major landmarks. Obviously, a higher number of landmarks may be used also in the computational scheme according to the present invention.

- Circular profiles

An alternative way for selecting the points in the neighborhood of a landmark is to sample points on a circle with center at the landmark and at least one radius. An example of a circular sampling is given in fig. 9. If the landmark is located at (x_0, y_0), the gray value function of the image is sampled at radius r_c at the points

$$(x_0 + r_c \cos\frac{2\pi}{n_c}(i-1), y_0 + r_c \sin\frac{2\pi}{n_c}(i-1)) \qquad i = 1,...,n_c .$$

The samples are put in a profile vector of length n_c. Suitable choices of n_c are 12, 8, 6, 4 and 3 (corresponding to 30, 45, 60, 90 and 120 degree angular spacing). Multiple circular sub-samplings at a set of radii may be considered simultaneously. Suitable radii r_c, expressed dimensionless as a fraction of the image size, are 0.05 and 0.025.

This scheme has the advantage over the linear profiles that it captures the 2D structure around the neighborhood. Specific

anatomical landmarks such as the costophrenic corner point or the intersection of the left heart shadow with the left diaphragm are better modeled by circular profiles than linear profiles for they have discontinuous first derivative of the lung field contour. Landmarks along the lung field contour in the vicinity of these specific anatomical landmarks may also ambiguously be represented by a linear profile if its length is too long, for then the profile may cross the juxtaposed lung field border.

The circular profile may be normalized such that the sum of absolute values of the elements in the profile is 1. The degree of sub-sampling may range from no sub-sampling (yielding a re-sampling at original resolution of the image) to very sub-sampled, e.g. only 4 points on the main axes in the case of a circular profile.

In three dimensions, a novel scheme to sample the neighborhood of a landmark in the volume is adopted. An appropriate choice is to sample on a sphere centered at the landmark with coordinates (x_0, y_0, z_0) and having radius r_c. Isotropic 3D sampling of the (x, y, z) coordinates on the sphere must be achieved, i.e. the point density on the spherical surface must be equal everywhere. When a spherical polar coordinate system is adopted, with φ denoting the angle w.r.t. the positive x axis in the xy plane, and θ the angle w.r.t. the z axis, and r_c the radius w.r.t. the origin centered at the landmark (fig. 15), the mapping to Cartesian coordinates is

$$\begin{aligned} x &= r_c \sin\theta \cos\varphi \\ y &= r_c \sin\theta \sin\varphi \\ z &= r_c \cos\theta \end{aligned} \quad .$$

φ (the longitude) can be sampled directly in a uniform way between $[0..2\pi]$ as in the 2D case, i.e. $\varphi = \frac{2\pi}{n_1}(i-1) \quad i = 1,...,n_1$. However, for equal spacing of θ (the latitude), the projections on the z axis

are more dense near the poles, hence the sampling in z can only be made uniform when $\cos\theta$ is sampled uniformly between $[-1..1]$ instead of θ being sampled uniformly between $\left[-\frac{\pi}{2},\frac{\pi}{2}\right]$. Hence $\theta = \arccos(\frac{j}{n_2})$ $j=-n_2,...,n_2$ is the sampling for θ. $t_1=\frac{2\pi}{n_1}$ and $t_2=\frac{1}{n_2}$ are the periods for the longitude resp. latitude.

A number of concentric spheres may be considered simultaneously to sample the neighborhood. The advantage of the spherical sampling is that the Euclidean distance, defined as $d_2((x_1,y_1,z_1),(x_2,y_2,z_2)) = ((x_2-x_1)^2+(y_2-y_1)^2+(z_2-z_1)^2)^{1/2}$, of all samples to the landmark's position is equal. Alternatively, the computationally less expensive city-block distance, defined as $d_1((x_1,y_1,z_1),(x_2,y_2,z_2)) = |x_2-x_1|+|y_2-y_1|+|z_2-z_1|$, may be used, producing samplings on a diamond shape centered on the landmark, with vertices on the axes; or the L^∞ norm or chess-board distance, defined as $d_\infty((x_1,y_1,z_1),(x_2,y_2,z_2)) = \max(|x_2-x_1|,|y_2-y_1|,|z_2-z_1|)$, producing cubes as iso-distance surfaces, with sides parallel to the axes. The latter scheme has the advantage that it may avoid interpolation in an isotropic volume, when the distance from the center is chosen as a multiple of the voxel size.

- Gray-level feature images

The image structure around a landmark is captured in a number N of mathematical features as outlined in the summary of the invention. The feature images are obtained by computing a number of moments of the local histogram of derivative images in a window with width α around each location x_0. In a preferred embodiment the first and second histogram moments, all derivatives up to second-order $(L, L_x, L_y, L_{xx}, L_{yy}, L_{xy})$ and 5 inner scales (σ=0.5,1,2,4,8 pixels) are computed, amounting to a total of $N=2\times6\times5=60$ feature images. The window scale to compute the histogram is in accordance to the

inner scale, i.e. $\alpha = 2\sigma$. The set of feature images can be regarded as a gray-level decomposition of the gray-level function in the neighborhood of a landmark.

- Gray-level cost

The Mahalanobis distance between a new profile $\mathbf{g}_{i,j}$ sampled in a new feature image j for a certain landmark i is computed as

$$h_{i,j}(\mathbf{p}) = \left(\mathbf{g}_{i,j}(\mathbf{p}) - \overline{\mathbf{g}}_{i,j}\right)^T S_{i,j}^{-1}\left(\mathbf{g}_{i,j}(\mathbf{p}) - \overline{\mathbf{g}}_{i,j}\right)$$

A smaller Mahalanobis distance means a larger probability that the profile $\mathbf{g}_{i,j}(\mathbf{p})$ originates from the Gaussian distribution with mean $\overline{\mathbf{g}}_{i,j}$ and covariance $S_{i,j}$. Hence, the Mahalanobis distance may be used as a gray-level cost measure, denoted as $h_{i,j}(\mathbf{p})$. This cost measure is a function of the location $\mathbf{p}$ in a given image. The location that most likely is the true location for the landmark i, in accordance to feature j, is the point $\mathbf{p}$ that minimizes $h_{i,j}(\mathbf{p})$. This gray-level cost can thus be defined for each feature j.

A total gray-level cost is obtained by combining all gray-level costs for all features j in a sum

$$h_i(\mathbf{p}) = \sum_{j=1}^{N}\left(\mathbf{g}_{i,j}(\mathbf{p}) - \overline{\mathbf{g}}_{i,j}\right)^T S_{i,j}^{-1}\left(\mathbf{g}_{i,j}(\mathbf{p}) - \overline{\mathbf{g}}_{i,j}\right)$$

reflecting the similarity between the gray-level pattern at $\mathbf{p}$ and the expected gray-level pattern at landmark i. The location most likely to coincide with landmark i, in accordance to the gray-level appearance, is the solution of the optimization problem

$$\mathbf{p}_o = \arg\min_{\mathbf{p}} h_i(\mathbf{p}) \ .$$

To construct the gray-level appearance model, the distributions $(\overline{\mathbf{g}}_{i,j}, S_{i,j})$ of the feature profiles $\mathbf{g}_{i,j}$ (for all landmarks $i=1,...,n$, for all feature images $j=1,...,N$ have to be estimated from the training images.

In three dimensions, the cost function at a certain 3D point $\mathbf{p}$ is likewise evaluated as the Mahalanobis distance $h_{i,j}(\mathbf{p})$ between a new 3D voxel-level profile $\mathbf{g}_{i,j}$ sampled spherically at the 3D point $\mathbf{p}$ in a new 3D feature image j for a certain 3D landmark i. A total gray-level cost is likewise obtained by combining all gray-level costs for all features j in a sum $h_i(\mathbf{p})$, reflecting the similarity in a spherically sampled neighborhood between the gray-level pattern at $\mathbf{p}$ and the expected gray-level pattern at landmark i. To search the most likely location for the 3D landmark i, an optimization problem must be solved.

To construct the gray-level appearance model in the 3D case, the distributions $(\overline{\mathbf{g}}_{i,j}, S_{i,j})$ of the spherically shaped feature image samplings $\mathbf{g}_{i,j}$ (for all landmarks $i=1,...,n$, for all feature images $j=1,...,N$) have to be estimated from the training volume images.

Shape model

Whereas a gray-level cost validates a gray-level pattern or a pattern of features, a shape cost is defined to validate a connection between two successive landmarks.

The curve outline, one for each of the lung fields, is described by n points (landmarks). These landmarks are manually determined in a set of s training images, yielding a sequence of points $(x_1,y_1)...(x_n,y_n)=(\mathbf{p}_1,...,\mathbf{p}_n)$. These coordinate tuples are subsequently positioned in a vector $\mathbf{x}=(x_1,y_1,...,x_n,y_n)^T$, representing the shape. Considering a pair $(\mathbf{p}_i,\mathbf{p}_{i+1})$ of estimated positions of

successive landmarks. A shape cost is assigned to the vector $\mathbf{v}_i = \mathbf{p}_{i+1} - \mathbf{p}_i$, reflecting the plausibility of $\mathbf{v}_i$ w.r.t. its probability distribution. The distributions of $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$, assumed to have normal distributions around their mean, are estimated from the training shapes. The mean and covariance of $\mathbf{v}_i$ are noted as $\overline{\mathbf{v}}_i$ and $S_{\mathbf{v}_i}$ respectively. An example of these vector distributions is given in fig. 10. The vector distributions can be regarded as a first-order (tangential) geometric decomposition of the shape represented by the set of landmarks.

A novel shape cost validating a connection vector $\mathbf{v}_i$ between two successive landmarks $\mathbf{p}_i, \mathbf{p}_{i+1}$ (a vector in the plane being fully characterized by its orientation and its length) is the Mahalanobis distance between $\mathbf{v}_i$ and its mean $\overline{\mathbf{v}}_i$:

$$f_i(\mathbf{v}_i) = \left(\mathbf{v}_i - \overline{\mathbf{v}}_i\right)^T S_{v_i}^{-1} \left(\mathbf{v}_i - \overline{\mathbf{v}}_i\right).$$

A connection, which is unlikely to happen, because it has large Mahalanobis distance, will get a high shape cost. On the other hand, a zero cost will be assigned to a connection that equals the expected value.

In three dimensions, each landmark i is a node or vertex in the point-based representation connected with edges to a number of vertices $v_{i,j}$ $j = 1..V_i$ corresponding to the neighboring landmark nodes on the surface or inside the surface. Each edge is a 3D vector characterized by its orientation in 3D space and a length, and the 3D distribution of each edge is derived from the training image set. In practice, it means that each edge of the polygons covering the surface is learned from the training examples. In particular, when the polygons are triangles, the triples of connection vectors of the three sides of all triangles are modeled. When also points inside the volume are included, and polyhedral cell are defined between the

ensemble of points on the surface and inside the volume, the distribution of the edges of the polyhedral decomposition are learned from the training set. In particular, when the polyhedrons are tetrahedrons (composed of four points and six connecting edges, forming four triangular faces), the spatial distribution of each of the six edges is modeled. The covariance matrix $S_{\mathbf{v}_i}$ effectively models the spatial correlation between the connection vectors, because it captures correlated variation that may exist between the positions of different anatomic landmarks. Alternatively, using the triangulated surface representation or the tetrahedral volume representation, the distribution of the normals of the planar patches formed by the triangles may be modeled also.

Segmentation System II (fig. 4)

Knowledge about the gray-level appearance of the object in the image and about the shape of the object to be segmented is acquired during the training phase. This knowledge is subsequently used to segment the object in a new image according to the following steps.

Step 1. Rectangular search grids construction.

A rectangular search grid for each landmark i is constructed to constrain the possible locations of the landmark. Each rectangular grid is characterized by the parameters of geometric extent and grid spacing.

The grid spacing r_g should be in accordance with the radius r_c of the circular profiles of the gray-level appearance model. A suitable relationship is $r_g = f_g r_c$ with r_c the radius and f_g a fixed fraction. A typical fraction for the grid spacing is $f_g = 0.5$.

The grid extent and grid position is landmark specific and is represented by $x_{\min}$, $x_{\max}$, $y_{\min}$ and $y_{\max}$. The true (unknown) landmark location (x^*, y^*) is deemed to lie between the lower and upper borders:

$$x_{\min} < x^* < x_{\max}$$
$$y_{\min} < y^* < y_{\max}$$.

These conditions can only be guaranteed if the grid spans the whole image area (i.e. $x_{\min} = y_{\min} = 0$ and $x_{\max} = y_{\max} = 1$). A more efficient approach is to constrain the search to a relevant part of the image. Assuming that the probability distribution $p_x(x)$ of the x-coordinate of the landmark is normal with mean $\bar{x}$ and standard deviation σ_x (estimated from the training shapes). The grid borders $x_{\min}$ and $x_{\max}$, jointly defining the grid extent in the x-direction, are chosen such that

$$\int_{x_{\min}}^{x_{\max}} p_x(x)dx = f_c$$

with f_c some fraction representing the probability that the condition $x_{\min} < x^* < x_{\max}$ is true. A suitable value for the fraction f_c is 0.995. If the requirement is imposed that the extent is symmetric around the mean $\bar{x}$, i.e. $x_{\max} - \bar{x} = \bar{x} - x_{\min}$, the upper border $x_{\max}$ is the value that satisfies

$$\frac{1}{\sigma_x\sqrt{2\pi}} \int_{\bar{x}}^{x_{\max}} \exp\left(-\frac{(x-\bar{x})^2}{2\sigma_x^2}\right) dx = \frac{1}{2} f_c .$$

The lower border is then $x_{\min} = 2\bar{x} - x_{\max}$. The borders for the y coordinate are obtained analogously.

In three dimensions, the search grid for each landmark i is a parallelepiped aligned with the coordinate axes. The grid spacing r_g should be in accordance with the radius r_c of the circular profiles of the gray-level appearance model. A suitable relationship is $r_g = f_g r_c$ with r_c the radius and f_g a fixed fraction. A typical fraction for the grid spacing is $f_g = 0.5$.

The grid extent and grid position is landmark specific and is represented by $x_{\min}$, $x_{\max}$, $y_{\min}$, $y_{\max}$, $z_{\min}$, and $z_{\max}$. The true (unknown)

landmark location (x^*, y^*, z^*) is deemed to lie between the lower and upper borders:

$$x_{\min} < x^* < x_{\max}$$
$$y_{\min} < y^* < y_{\max} \ .$$
$$z_{\min} < z^* < z_{\max}$$

These conditions can only be guaranteed if the grid spans the whole image volume (i.e. $x_{\min} = y_{\min} = z_{\min} = 0$ and $x_{\max} = y_{\max} = z_{\max} = 1$. A more efficient approach is to constrain the search to a relevant part of the image. Similar to the 2D approach, the grid extent is derived from the training image set so as to cover a fraction f_c of the integrated probability distribution function of landmark positions in the 3D volumes. A suitable value for the fraction f_c is 0.995.

In contrast to the linear profile sampling, that restricts the solution for a landmark to lay on the 1D perpendicular to the current segmentation outline, the search grid enlarges the solution space to lie in a 3D neighborhood around the current landmark.

Step 2. Gray-level cost matrix construction

The gray-level appearance model is used to find proper locations for each landmark. The m best locations of landmark i, in accordance with the gray level appearance model, are selected as follows.

- First, a rectangular grid covering a relevant part of the image area around the expected location of each landmark i is defined, according to the previous step.
- Secondly, for each grid point, the total gray-level cost $h_i(\mathbf{p})$ is computed. The points $\mathbf{p}_{i,1}, \mathbf{p}_{i,2}, ..., \mathbf{p}_{i,m}$ corresponding to the m lowest total gray-level costs are selected.
- Thirdly, a gray-level cost matrix C is constructed so that each row i contains the costs of the m most likely locations of landmark i:

$$C = \begin{bmatrix} h_1(p_{1,1}) & \cdots & h_1(p_{1,k}) & \cdots & h_1(p_{1,m}) \\ \vdots & & \vdots & & \vdots \\ h_i(p_{i,1}) & \cdots & h_i(p_{i,k}) & \cdots & h_i(p_{i,m}) \\ \vdots & & \vdots & & \vdots \\ h_n(p_{n,1}) & \cdots & h_n(p_{n,k}) & \cdots & h_n(p_{n,m}) \end{bmatrix}.$$

A typical number m of best points per landmark ranges from 10 to 40.

As the m of best points per landmark are selected independently from the set of m of best points for a neighboring landmark, the situation may arise that one or more of these points are nearer to a non-neighboring landmark. These points will likely be neglected in the final contour segmentation by accounting for the shape cost in the next step.

In three dimensions, the equivalent of a gray-level cost matrix is the gray-level cost graph with each node having a multiplicity representing the number of m lowest total gray-level cost points for the landmark. As in the 2D case, the gray-level appearance model is used to find proper locations for each landmark. The m best locations of landmark i, in accordance with the gray level appearance model, are selected as follows.

- First, a 3D grid covering a relevant part of the image area around the expected location of each landmark i is defined, according to the previous step.
- Secondly, for each 3D grid point, the total gray-level cost $h_i(\mathbf{p})$ is computed, using the spherically sampled neighborhood. The points $\mathbf{p}_{i,1}, \mathbf{p}_{i,2}, \ldots, \mathbf{p}_{i,m}$ corresponding to the m lowest total gray-level costs are selected.
- Thirdly, a gray-level cost graph G is constructed, each node i having a multiplicity m, to contain the costs of the m most likely 3D locations of landmark i. The graph is constructed

starting from a node of the surface triangulation and, following the edges, linking successively all nearest neighbors on the surface. In each node, the m lowest total gray-level cost points of the previous step are stored. A typical number m of best points per 3D landmark ranges from 10 to 40.

Because the only criterion is gray-level cost, this step may lead to the situation that a node comprises points in its set of m most likely points that are nearer in 3D space to non-neighboring nodes of node i, especially when the landmarks are densely sampled over the surface. These points will likely be neglected in the final surface segmentation by accounting for the shape cost in the next step, i.e. they will represent unlikely connections.

Step 3. Minimal cost path construction

Determining the contour that segments the object is reduced to finding a path from top to bottom through the matrix C by selecting one element per row. Denoting the index of the selected element in row i as k_i, the curve becomes the point sequence $\{\mathbf{p}_{1,k_1},\mathbf{p}_{2,k_2},...,\mathbf{p}_{n,k_n}\}$. The optimal path $(k_1^*,k_2^*,...,k_n^*)$ is the path that minimizes a cost function $J(k_1,...,k_n)$:

$$(k_1^*,k_2^*,...,k_n^*)=\arg\min_{k_1,...,k_n} J(k_1,...,k_n).$$

The models introduced above admit a number of cost measures. Considering two successive landmarks p_{i,k_i} and $p_{i+1,k_{i+1}}$, a cost component according to the gray-level appearance model and a cost component according to the shape model are:

- The gray level costs $h_i(p_{i,k_i})$ and $h_{i+1}(p_{i+1,k_{i+1}})$ corresponding to the landmarks p_{i,k_i} and $p_{i+1,k_{i+1}}$;
- The shape cost $f_i(p_{i+1,k_{i+1}}-p_{i,k_i})$, which validates the plausibility of the connection from p_{i,k_i} to $p_{i+1,k_{i+1}}$.

Both cost types can be incorporated in the cost function $J(k_1,...,k_n)$ by means of a weighted sum of an overall gray-level cost and an overall shape cost. The overall gray-level cost with weight factor γ_1 is the sum of the landmark individual gray-level costs $h_i(p_{i,k_i}), i=1,...,n$. The overall shape cost with weight factor γ_2 is the sum of the shape costs $f_i(p_{i+1,k_{i+1}} - p_{i,k_i}), i=1,...,n$. Hence the cost function $J(k_1,...,k_n)$ becomes

$$J(k_1,...,k_n) = \gamma_1 \sum_{i=1}^{n} h_i(p_{i,k_i}) + \gamma_2 \sum_{i=1}^{n-1} f_i(p_{i+1,k_{i+1}} - p_{i,k_i}) + \gamma_2 f_n(p_{1,k_1} - p_{n,k_n}).$$

The optimal path $\left(k_1^*, k_2^*, ..., k_n^*\right)$ that minimizes $J(k_1,...,k_n)$ is called a Minimal Cost Path. The optimal path is computed using dynamic programming techniques known in the prior art, such as introduced in G. Behiels et al., Evaluation of image features and search strategies for segmentation of bone structures using active shape models, Medical Image Analysis, 6(1):47-62, 2002. A typical weight factor for the gray-level cost is $\gamma_1 = 1$, and for the shape cost $\gamma_2 = 0.25$.

Fig. 13 shows the automatic segmentation of three thorax images according to the system outlined before, and compares it to the manual segmentation performed by an experienced radiologist. The 14 auto-placed landmarks on each lung field are shown, along with a continuous spline interpolation through them.

In three dimensions, a minimal cost surface construction to determine the optimal segmented surface may likewise use the combination of a gray-level cost and a shape cost. As in the previous module using linear samplings, the path through the points of the surface graph may be in the form of a traversal of a Minimum Spanning Tree (MST) or a depth first tree.

Determining the surface that segments the object is reduced to finding one of the m points per landmark node on the surface.

Denoting the index of the selected element in node i as k_i, the surface becomes the point sequence $\{\mathbf{p}_{1,k_1},\mathbf{p}_{2,k_2},...,\mathbf{p}_{n,k_n}\}$ of the traversed tree. The optimal surface $(k_1^*,k_2^*,...,k_n^*)$ is the point set that minimizes a cost function $J(k_1,...,k_n)$:

$$(k_1^*,k_2^*,...,k_n^*)=\arg\min_{k_1,...,k_n} J(k_1,...,k_n).$$

The models introduced above admit a number of cost measures. Considering two neighboring landmarks p_{i,k_i} and $p_{i+1,k_{i+1}}$ connected by an arc, a cost component according to the gray-level appearance model and a cost component according to the shape model are:

- The gray level costs $h_i(p_{i,k_i})$ and $h_{i+1}(p_{i+1,k_{i+1}})$ corresponding to the neighboring landmarks p_{i,k_i} and $p_{i+1,k_{i+1}}$ on the surface (or, additionally, inside the surface) connected by an arc;
- The shape cost $f_i(p_{i+1,k_{i+1}}-p_{i,k_i})$, which validates the plausibility of the connection vector in 3D space from p_{i,k_i} to $p_{i+1,k_{i+1}}$, the vector being fully characterized by its orientation and length.

Both cost types can be incorporated in the cost function $J(k_1,...,k_n)$ by means of a weighted sum of an overall gray-level cost and an overall shape cost. The overall gray-level cost with weight factor γ_1 is the sum of the landmark individual gray-level costs $h_i(p_{i,k_i}),i=1,...,n$. The overall shape cost with weight factor γ_2 is the sum of the shape costs $f_i(p_{i+1,k_{i+1}}-p_{i,k_i}),i=1,...,n$. Hence the cost function $J(k_1,...,k_n)$ becomes

$$J(k_1,...,k_n)=\gamma_1\sum_{i=1}^{n}h_i(p_{i,k_i})+\gamma_2\sum_{i=1}^{n-1}f_i(p_{i+1,k_{i+1}}-p_{i,k_i})+\gamma_2 f_n(p_{1,k_1}-p_{n,k_n}).$$

The sequence $p_1,...,p_n$ represents the chain of visited nodes as the tree is traversed. The optimal surface $(k_1^*,k_2^*,...,k_n^*)$ that minimizes $J(k_1,...,k_n)$ is called a Minimal Cost Surface.

Computational speed-up refinements

The segmentation method consisting of steps 1-3 can be applied to segment only part of the outline of one or more anatomic entities. To this purpose the rectangular search grid is only constructed for a subset of consecutive landmarks. The gray level cost matrix will comprise a number of rows equal to the number of retained landmarks. The minimum cost path construction will minimize a cost function that only comprises a gray level term and a shape term for the retained landmarks. The resulting segmentation traces only that outline part that is covered by the retained landmarks. The obvious advantage of partial segmentation is computational speed when one is only interested in part of the anatomic entity outline, or when only some specific landmarks need to be located to obtain specific measurements.

Another algorithmic refinement to speed up the computation of the segmentation is to use a coarse-to-fine multi-resolution approach. The gray-level appearance and shape models are constructed on a multi-resolution representation for each landmark, introducing the spatial scale as a new dimension. Two alternatives are implemented to this purpose, both applying the segmentation concept over successively finer resolution levels.

In one alternative, during segmentation, the number of points in a profile, the number of points in the search grid and the number of landmark points along the contour are decreased in the original image. The positional precision of the landmarks is then increased by successively increasing the resolution of the search grid in the vicinity of the previous solution (the total number of considered points in the profile and search grid may thus be held constant at each iteration), and applying the finer resolution gray-level and shape models.

In another alternative, the multi-resolution approach may be implemented using a number of levels of a Gaussian and Laplacian pyramid and their associated derivative feature images. The initial position of the landmarks is determined on a coarse level, and is

tracked to its final position on the finest level using the intermediate positions of the previous level as the starting position for the next level. Because the coarse resolution images contain fewer details, the search space will likely contain less false minima which will enable faster optimization at a given level and faster locating the landmarks towards their final position.

Training the segmentation models

A number of thoracic images is collected from an image repository, and presented sequentially in a Graphical User Interface. The following steps are performed to build the segmentation models:

- An experienced user manually segments the lung fields in each displayed image, by indicating -using left mouse clicks- a number of points along the contour that delineates the lung field. These points need not be spaced equidistantly along the lung field outline; the only requirement is that the point set collectively approximates the lung field to a sufficiently high degree. To asses the anatomic fit, a spline curve is continually interpolated through the manually positioned points determined so far, until the curve is closed from the last point to the first point by a right mouse click. Manual segmentation adjustments can be made by dragging an individual point towards a new position. The resulting outline is again assessed with respect to anatomical correctness.
- Next, a number of landmarks will be auto-placed on the manual segmented lung field contour. In order to achieve that identical landmarks on all images of the training set map on each other, the user positions a few number of easily discernible landmarks, the other landmarks are obtained by equidistantly placing points on the lung field contour. In the case of lung field outlines, a number of easily discernable landmarks are the topmost lung field point, the point denoting the costophrenic angle and the junction of heart shadow and hemi-diaphragm, a sub-total of three. Next a number of total landmarks are chosen, a suitable choice ranging from 14 to 40 for example. In the 14 landmark case, the points $\mathbf{p}_1$, $\mathbf{p}_7$ and $\mathbf{p}_9$ represent the three fixed landmarks, five points are evenly

distributed between $\mathbf{p}_1$ and $\mathbf{p}_7$, one point between $\mathbf{p}_7$ and $\mathbf{p}_9$, and another five points between $\mathbf{p}_9$ and $\mathbf{p}_1$. This step is performed separately on left and right lung fields.

- Next, parameters pertaining to the training phase are asked: (a) image size for training (and segmentation), a typical value is 256 or 512; (b) fraction of the shape variance to be explained by the Principal Components Analysis (PCA), a typical value being 0.9999; (c) number of points in the profile, either linear or circular, a typical value for the circular profile being 12, 8, 6, 4 or 3; (d) radius of the circular profile as a fraction of the image dimension, typical values being 0.04, 0.03, 0.02.

- Training of the gray-level appearance model: (a) decomposition of the gray-level function around the landmarks, i.c. computation of N feature images e.g. the local histogram moments of gray-value derivatives as outlined before, and (b) collection of the gray value distribution of the linear or circular profiles at the landmarks $\mathbf{p}_i$, i.c. computation of $\overline{\mathbf{g}}_{i,j}, i=1...n, j=1...N$ and the covariance matrices $S_{i,j}, i=1...n, j=1...N$ (i.e. for each landmark i in feature image j). This step generates the statistical gray-level appearance knowledge.

- Training of the shape model: (a) computation of geometric decomposition of the zeroth order (positional) information contained in the landmarks, i.c. computation of the mean shape $\overline{\mathbf{x}}$ and the t principal modes of variation (eigenshapes) arranged columnwise in a matrix $\mathbf{b}$; (b) computation of the vector distributions (first order tangential information) contained in the connection sequence of landmarks, i.c. computation of the mean $\overline{\mathbf{v}}_i$ and covariance $S_{\mathbf{v}_i}$ of $\mathbf{v}_i$. This step generates the statistical shape knowledge.

- Storage of the statistical gray-level appearance and shape knowledge, e.g. to be used to segment the lung fields in a new image according to the model-based segmentation sub-systems given above.

Application of the model-based segmentation in 2D images to the computation of the cardio-thoracic ratio (CTR)

An application of the automated lung field segmentation is the computation of the Cardiothoracic Ratio (CTR). The CTR (fig. 6) is defined as the ratio of the transverse diameter of the heart to the internal diameter of the thorax (ID):

$$CTR = \frac{MLD + MRD}{ID}$$

with MLD the maximum transverse diameter on the left side of the heart and MRD the maximum transverse diameter on the right side of the heart. This index is an important clinical parameter, which varies for an adult between 39% and 50% with an average of about 45%. A cardiothoracic index higher than 50% is considered abnormal. Possible causes are cardiac failure, pericardial effusion and left or right ventricular hypertrophy. It is possible to compute the cardiothoracic ratio automatically, using the automatic lung field segmentation as disclosed in the present invention.

Referring to fig. 8, the characteristic point defining MRD is obtained by selecting the Cartesian leftmost point on the fitted contour segment between landmarks p1 and p9 of the right lung segmentation; the characteristic point defining MLD is obtained by selecting the Cartesian rightmost point on the fitted contour segment between landmarks p1 and p7 of the left lung segmentation. The sum MLD+MRD is obtained by subtracting the column coordinates of these characteristic points and taking the absolute value. Similarly, the ID is obtained by selecting the Cartesian leftmost point on the fitted contour segment between landmarks p1 and p7 of the right lung segmentation and the Cartesian rightmost point on the fitted contour segment between landmarks p1 and p9 of the left lung segmentation, subtracting the column coordinates of these characteristic points and taking the absolute value.

Fig. 14 shows the computation of the characteristic points on the segmented lung fields, and the computation of the cardiothoracic ratio using (a) the manual lung field segmentations resulting in a

$CTR_{man} = 0.47$; (b) using the automatically derived lung field segmentations according to the present invention resulting in a $CTR_{automatic} = 0.45$.

Application of the model-based segmentation and measurement system to other body parts

The spatial extent of the search grid for each of the landmarks in the lung field segmentation is derived on the basis of the positions of all similar landmarks in a set of training thorax images. The concept of using a search grid for constraining the candidate positions for a given landmark can be extended for other body parts with that may have wider positional, rotational and size variation in the image than that of a lung field in a thorax image, and that do occupy the full image area as opposed to lung fields in a thorax image.

To allow for such positional, rotational and size variation of the body part, the concept of anchor point mapping of the search grids using methods disclosed in European patent application 04076454, entitled "Method for automatically mapping of geometric objects in digital medical images", may be applied in conjunction with the present invention. The refined method then becomes particular interesting to segment bony body parts in orthopedic radiographs because they are characterized by a relatively constant shape in the image, and they can be anchored to well manifested bony landmarks in the image.

For example, in a pelvis, hip or full leg examination, the femoral outlines can be anchored to well known landmarks such as the tip of the greater trochanter, the knee center and the femoral head center. These anchor points are used to establish an affine transformation between model anchor points and the associated anchor points selected in the actual image. Because a search grid consists of a collection of candidate points arranged in a rectangular lattice around a landmark point on the segmentation model contour in the model image, each of the constituent candidate points is mapped in turn in the actual image by applying the transformation to the model

points' coordinates. In this way, the search grid for a given landmark point is reconstructed around the most likely position of the point. The optimization process then proceeds in the manner disclosed above by optimizing a combined gray value and shape cost for a path through a selected combination of candidate locations, one on each search grid. The path with most optimal cost is the final segmentation of the body part. Examples of other bones are hand and upper extremity bones, foot and lower extremity bones, pelvis and spinal vertebrae and structures.

Other body parts that are amenable to this type of landmark-based segmentation are soft-tissue organs in 3D CT of MR images. On the one hand, a slice-by-slice based approach may be taken here, that determines the segmented contour of the anatomy on each slice, and that combines the result of all slices into a 3D segmented surface. On the other hand, a fully 3D approach, extending the current concept on a 3D volume, may be adopted, that builds and applies models for 3D landmarks confined within 3D search grids. Kidneys, lungs, heart, liver, stomach, spleen, prostate and brain structures are examples of organs to which the method of the present invention is applicable.

In all application examples mentioned, measurement points may be based directly or indirectly on the position of or the landmarks on the segmentation or on a combination of several landmarks (such as a midpoint of two juxtaposed points on either side of the cortex of an elongated bone, a pair of which determines the bone's anatomical axis), and the landmark points may be subsequently fed into a measurement system such as disclosed in EP A 1349098.

Application of the model-based segmentation in 3D images to the computation of total lung volumes (TLV)

The approach for segmenting objects in 2D medical images can be extended as previously to segment objects embedded in 3D volumetric image data such as CT or MR. The resulting segmentations may be used to derive quantitative measurements to aid in the diagnosis. These measurements include size measurements such as organ volumes or the volumes of pathological tissues. Quantitative measurements may be

used to compare an individual patient to his prior measurements to track disease progression or therapy evolution, or compare a patient to other patients, or to reference values. A measurement quantity frequently derived from volume images is the volume V of an object or organ, represented by the voxels enclosed by the segmented surface S, and computed as

$$V = \sum_{v \in V} \Delta x \Delta y \Delta z ,$$

where Δx and Δy represent the in-slice voxel dimensions and Δz represents the between-slice spacing.

For example, the total lung volume (TLV) can be calculated on the basis of 3D lung field segmentation, and compared against the prior TLV to assess change. Comparison against reference values for the TLV according to age and gender may be used to assess whether the lung volume is indicative of emphysema. Obviously, the TLV computed on the basis of CT lung field segmentations can be divided in a left and right lung volume, which discrimination cannot be obtained by respiratory lung function tests.

The method according to this invention to construct models of anatomic objects and segment these anatomic objects in medical images can be applied to other organs. Examples are the hip or knee outlines on 2D radiographs and bones, brain, thoracic and abdominal structures on CT or MR images. Subsequent measurements can be based on the geometric information that is contained in the segmented surface of the object(s) or, more directly, be based on the positions of the defining landmarks lying on or inside the resulting segmented surface.

■

[CLAIMS]

1. A method of generating a gray value model associated with an anatomic entity in a three dimensional digital medical image characterised by the steps of

- sampling a manually segmented surface of said anatomic entity at a number of landmark points;
- forming a polygonal mesh by connecting groups of neighboring landmark points on the surface in a polygon;
- forming a graph data structure from the polygonal mesh, each node (vertex) of the graph representing a landmark and each edge (arc) of the graph representing a connection between landmarks;
- sampling a number of points in a neighborhood of each of said landmark points;
- for each landmark point arranging sampled points in a neighborhood around a landmark point in a profile;
- computing at least one feature image;
- computing a mean profile for each landmark point and for each feature image;
- computing a covariance matrix of the profiles for each landmark point and each feature image;
- identifying said graph, mean profiles and covariance matrices as gray value model of said anatomic entity.

2. A method of generating a gray value model associated with an anatomic entity in a three-dimensional digital medical image characterized by the steps of

- sampling a manually segmented volume of said anatomic entity at a number of landmark points;
- forming a polyhedral mesh by connecting groups of neighboring landmark points in the volume in a polyhedron;
- forming a graph data structure from the polyhedral mesh, each node (vertex) of the graph representing a landmark and

each edge (arc) of the graph representing a connection between landmarks;

- sampling a number of points in a neighborhood of each of said landmark points;
- for each landmark point arranging sampled points in a neighborhood around a landmark point in a profile;
- computing at least one feature image;
- computing a mean profile for each landmark point and for each feature image;
- computing a covariance matrix of the profiles for each landmark point and each feature image;
- identifying said graph, mean profiles and covariance matrices as gray value model of said anatomic entity.

3. A method according to claim 1 or claim 2 wherein a feature image is a derivative image at a predefined scale σ or an image representation comprising mathematical moments of the local histogram of said three-dimensional image or derivative image in a window with width α around the location of a landmark.

4. A method of generating a geometric model of an anatomic entity in a three dimensional digital medical image characterised by the steps of

- sampling a manually segmented surface of said anatomic entity at a number of landmark points;
- forming a polygonal mesh by connecting groups of neighboring landmark points on the surface in a polygon;
- forming a graph data structure from the polygonal mesh, each node (vertex) of the graph representing a landmark and each edge (arc) of the graph representing a connection between landmarks;
- computing connection vectors or connection vector differences between pairs of landmark points in the graph of the polygonal mesh;

- computing a mean connection vector or means vector difference between pairs of landmark points of the graph;
- computing a covariance matrix of connection vectors for each connection in the graph;
- identifying said graph, mean connection vectors or mean connection vector differences and covariance matrices as a geometric model of said anatomic entity.

5. A method of generating a geometric model of an anatomic entity in a three-dimensional digital medical image characterized by the steps of

- sampling a manually segmented volume of said anatomic entity at a number of landmark points;
- forming a polyhedral mesh by connecting groups of neighboring landmark points on the surface in a polyhedron;
- forming a graph data structure from the polyhedral mesh, each node (vertex) of the graph representing a landmark and each edge (arc) of the graph representing a connection between landmarks;
- computing connection vectors between pairs of landmark points of the graph;
- computing a mean connection vector between pairs of landmark points of the graph;
- computing a covariance matrix of connection vectors for each connection in the graph;
- identifying said graph, mean connection vectors and covariance matrices as a geometric model of said anatomic entity.

6. A method according to claim 1 where the polygonal mesh is a surface triangulation.

7. A method according to claim 2 where the polyhedral mesh is a volume tetrahedralization.

8. A method according to claim 1, 2, 4 or 5 wherein said gray value model and said geometric model are constructed from and applied to a multi-resolution representation of the image.

9. A computer program product adapted to carry out the method of any of the preceding claims when run on a computer.

10. A computer readable medium comprising computer executable program code adapted to carry out the steps of any of claims 1, 2, 4 or 5.

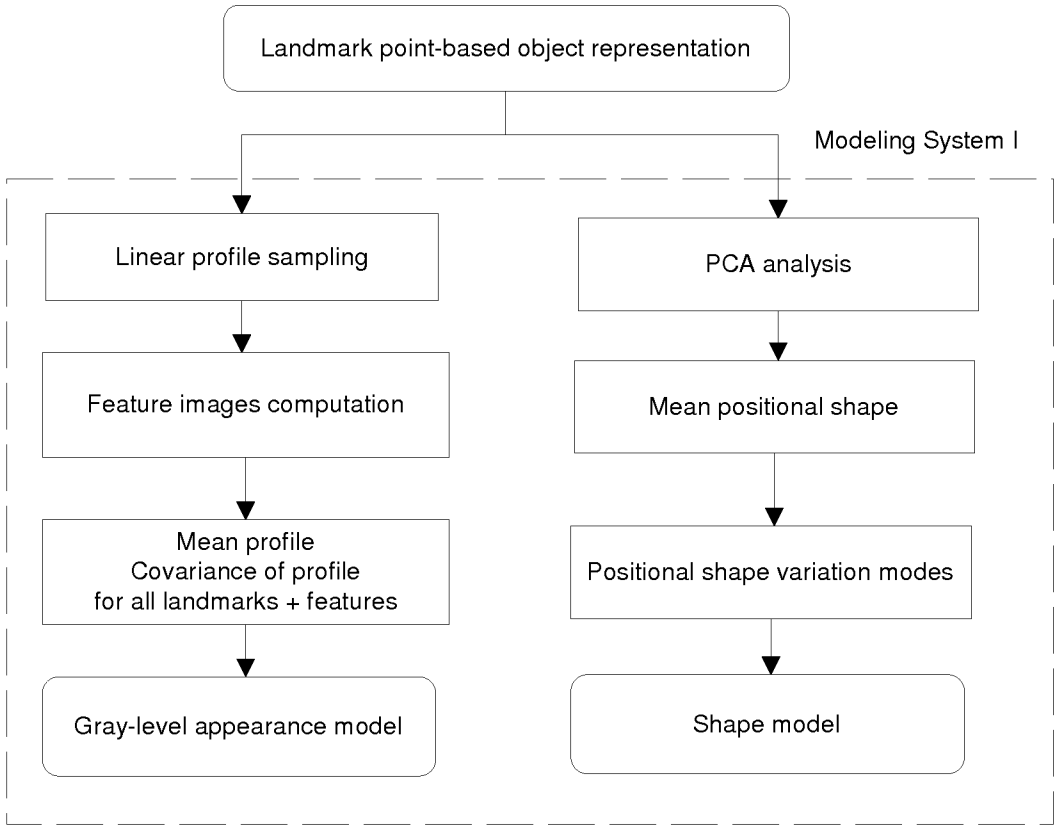

Landmark point-based object representation
Modeling System I
Linear profile sampling
Feature images computation
Mean profile
Covariance of profile
for all landmarks + features
Gray-level appearance model
PCA analysis
Mean positional shape
Positional shape variation modes
Shape model

FIG. 1

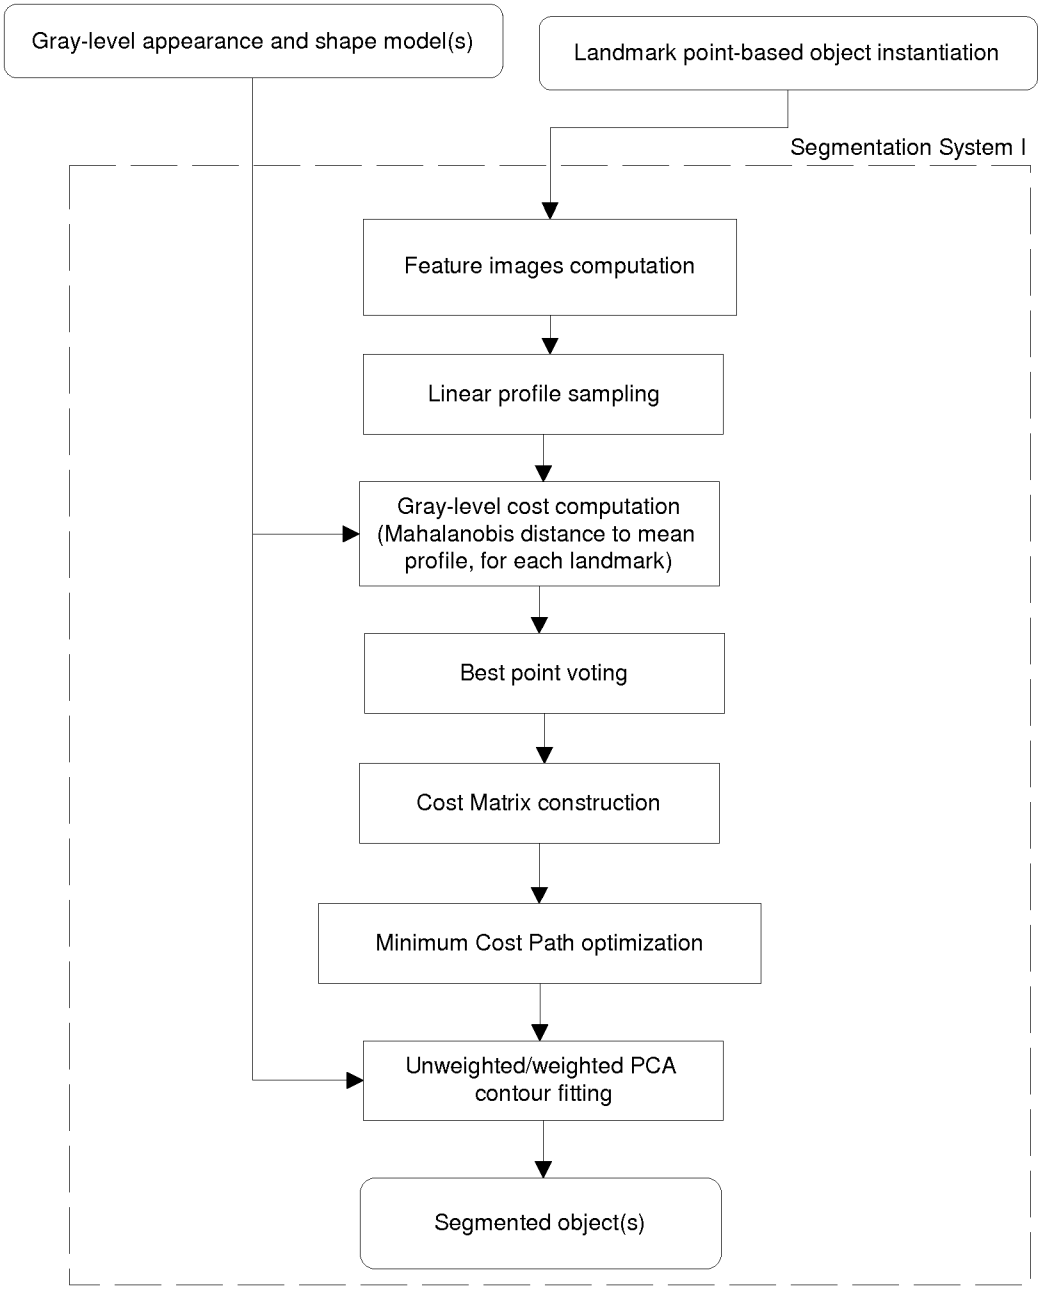

Gray-level appearance and shape model(s)
Landmark point-based object instantiation
Segmentation System I
Feature images computation
Linear profile sampling
Gray-level cost computation (Mahalanobis distance to mean profile, for each landmark)
Best point voting
Cost Matrix construction
Minimum Cost Path optimization
Unweighted/weighted PCA contour fitting
Segmented object(s)

FIG. 2

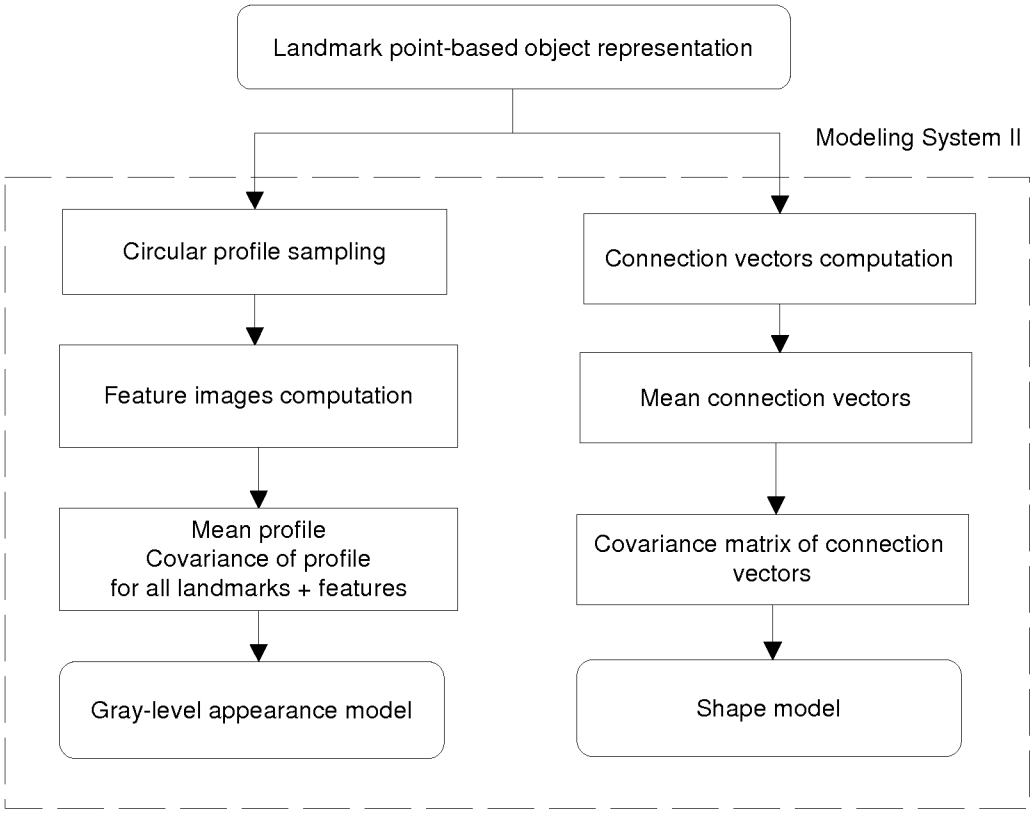

Landmark point-based object representation
Modeling System II
Circular profile sampling
Feature images computation
Mean profile
Covariance of profile
for all landmarks + features
Gray-level appearance model
Connection vectors computation
Mean connection vectors
Covariance matrix of connection vectors
Shape model

FIG. 3

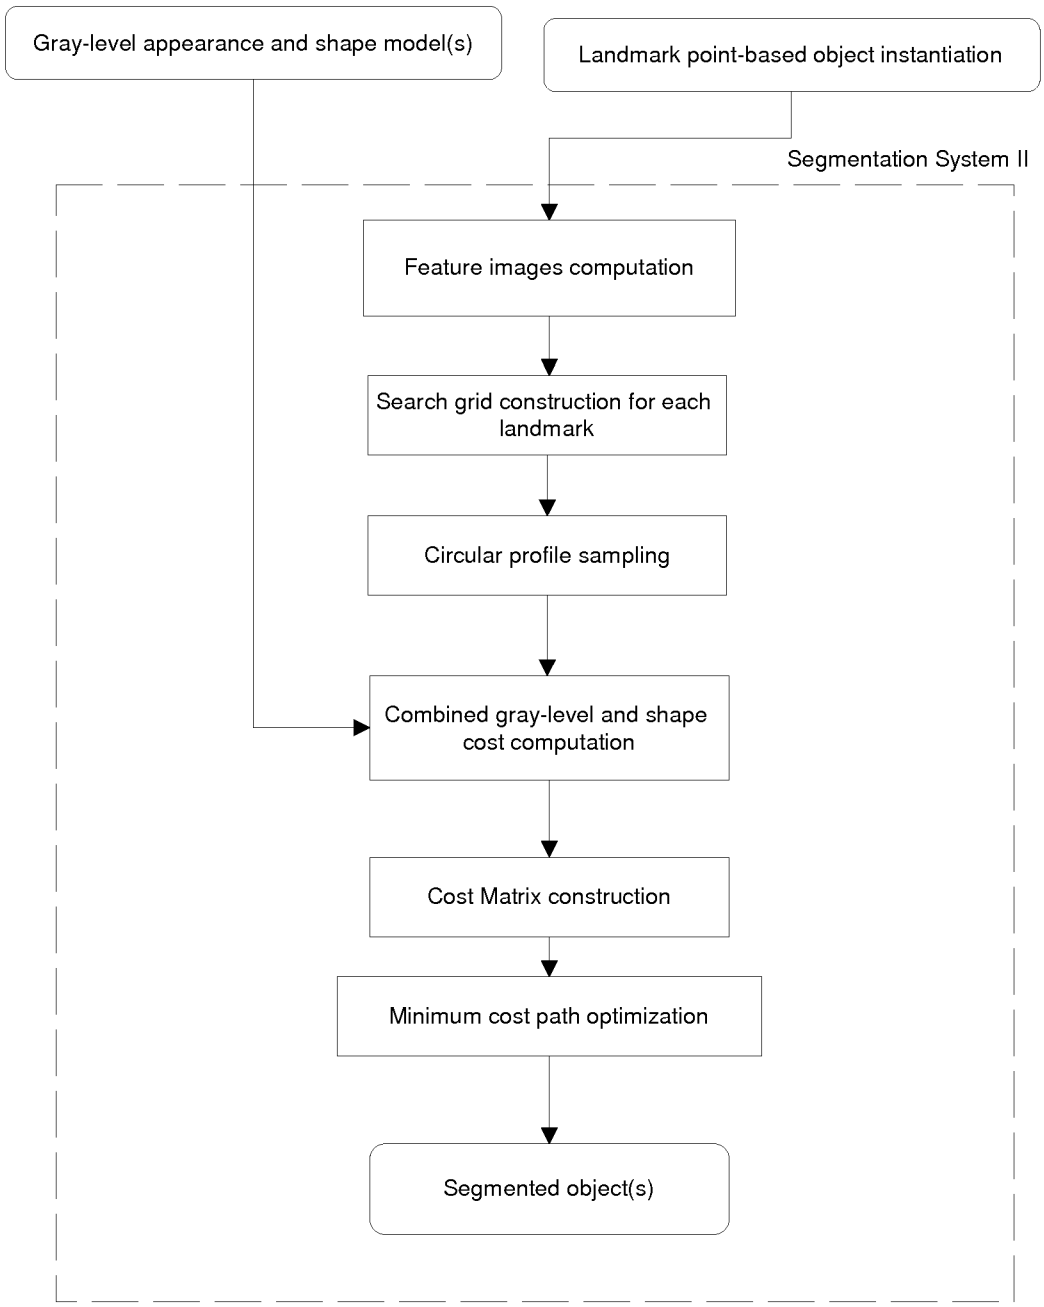

Gray-level appearance and shape model(s)
Landmark point-based object instantiation
Segmentation System II
Feature images computation
Search grid construction for each landmark
Circular profile sampling
Combined gray-level and shape cost computation
Cost Matrix construction
Minimum cost path optimization
Segmented object(s)

FIG. 4

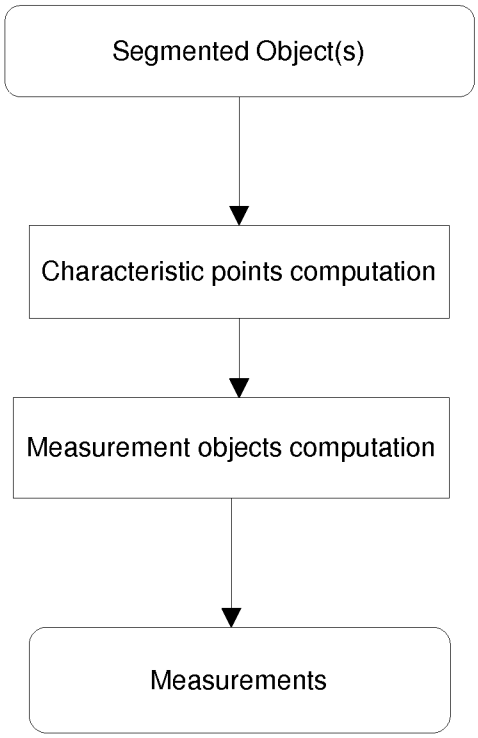

Segmented Object(s)
Characteristic points computation
Measurement objects computation
Measurements

FIG. 5

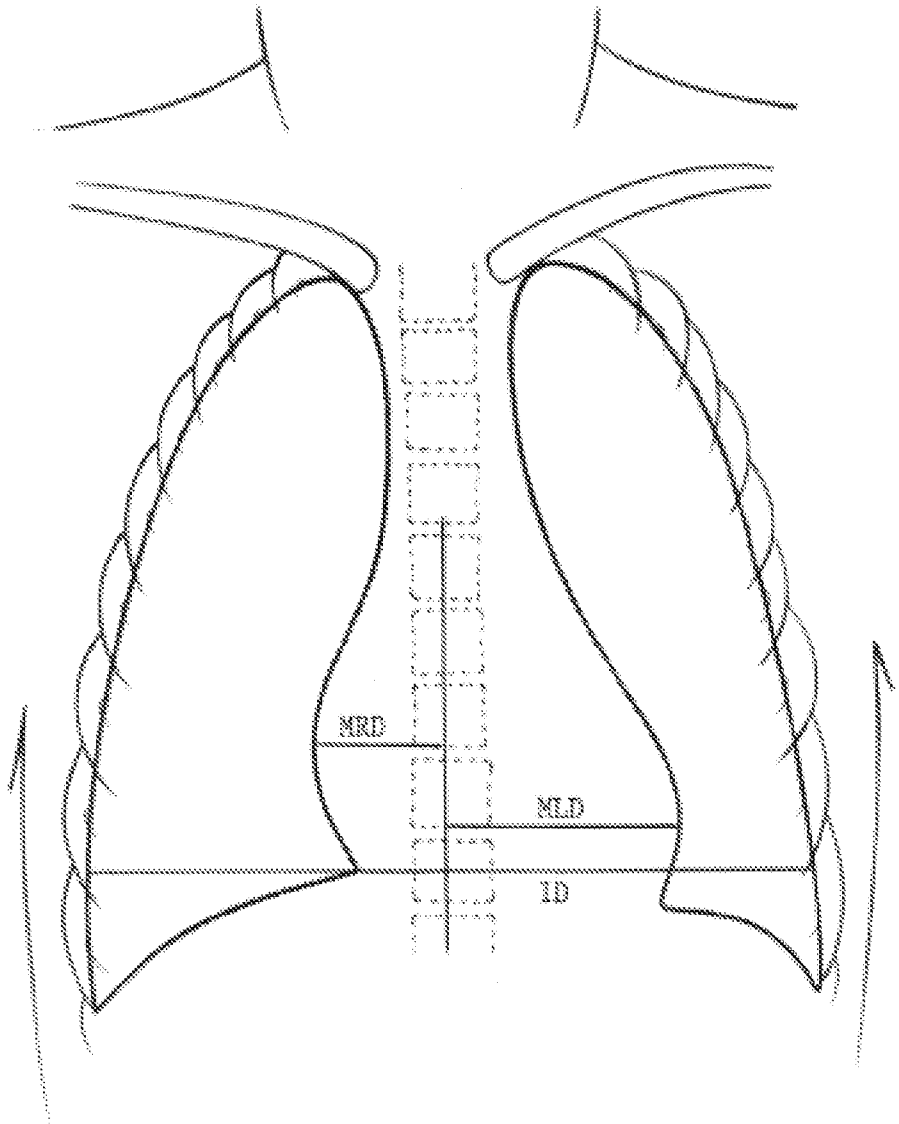


FIG. 6

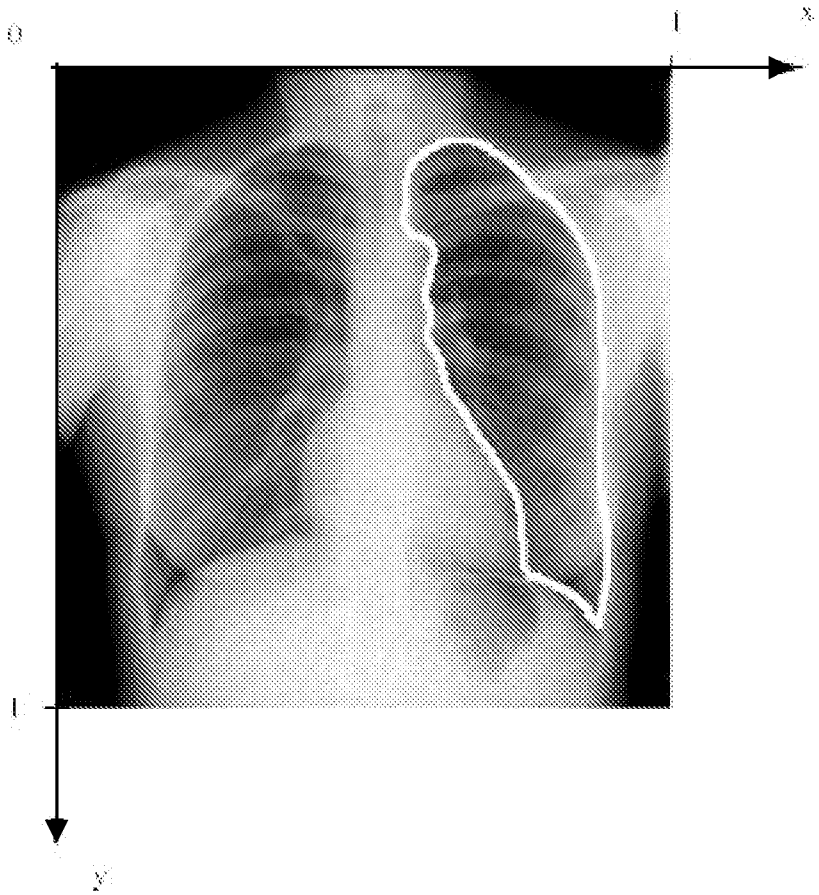


FIG. 7 (a)

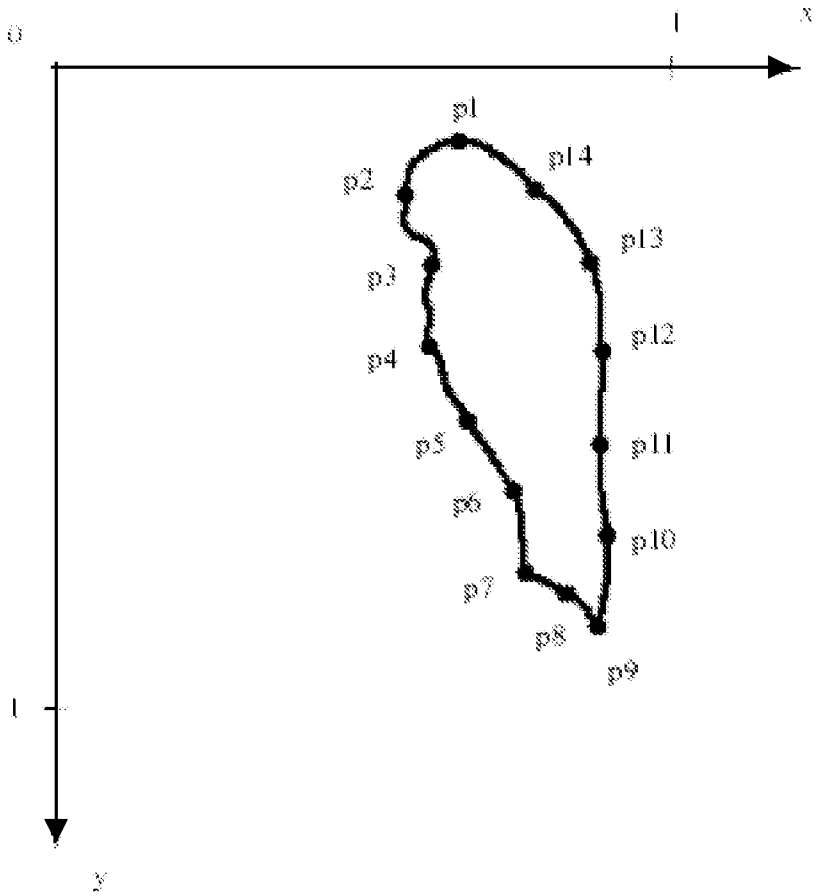


FIG. 7 (b)

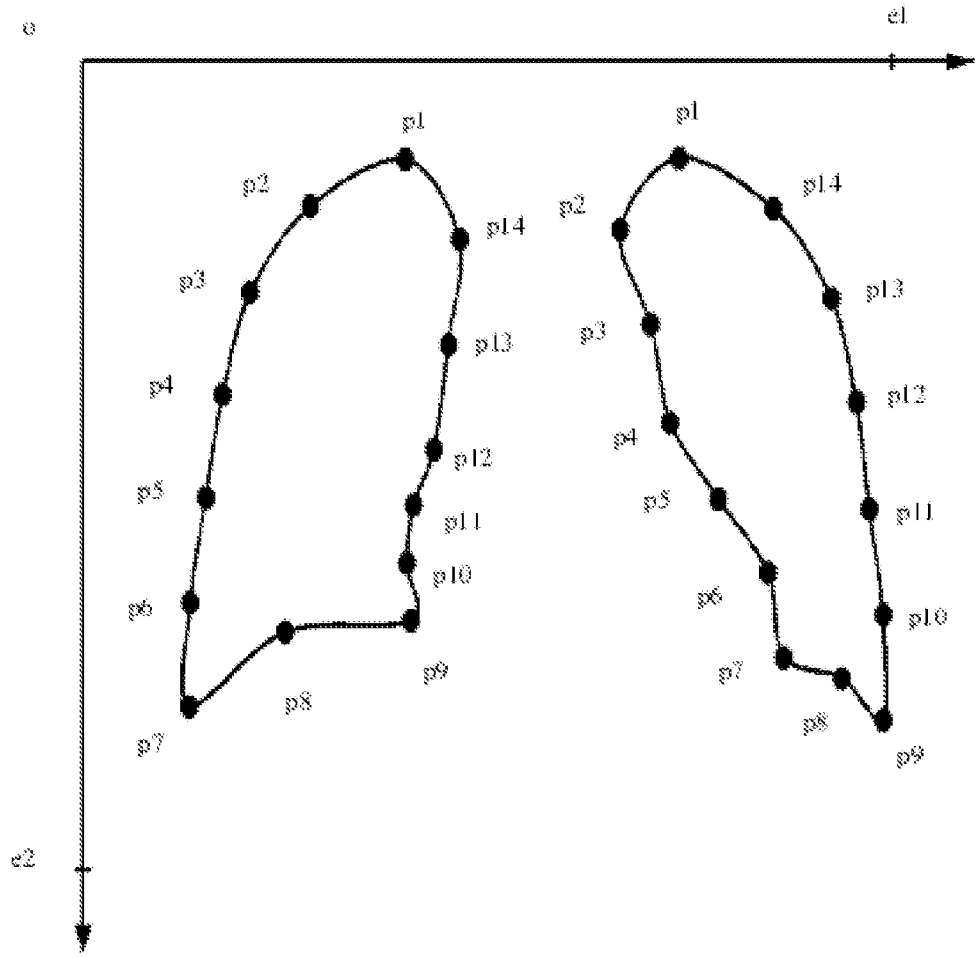

o
e1
e2
p1
p2
p3
p4
p5
p6
p7
p8
p9
p10
p11
p12
p13
p14
p1
p2
p3
p4
p5
p6
p7
p8
p9
p10
p11
p12
p13
p14

FIG. 8

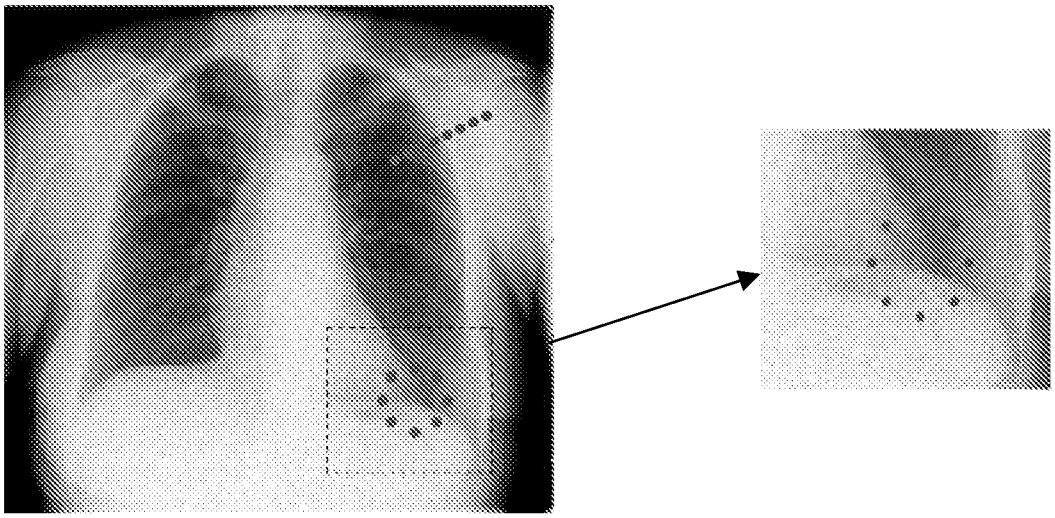
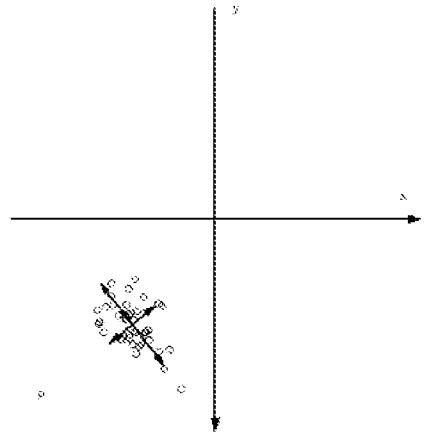
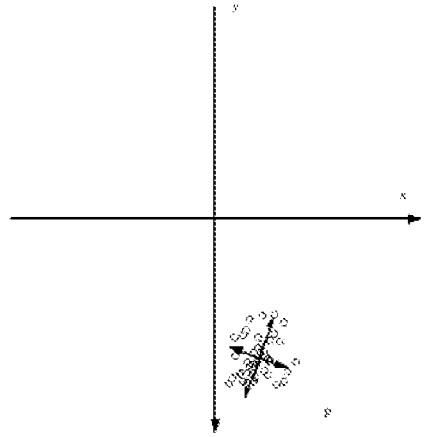


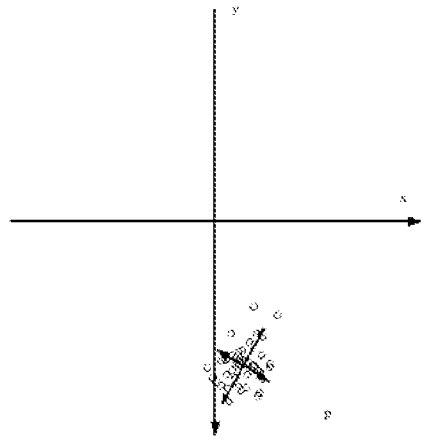
FIG. 9



(a)



(b)



(c)

FIG. 10

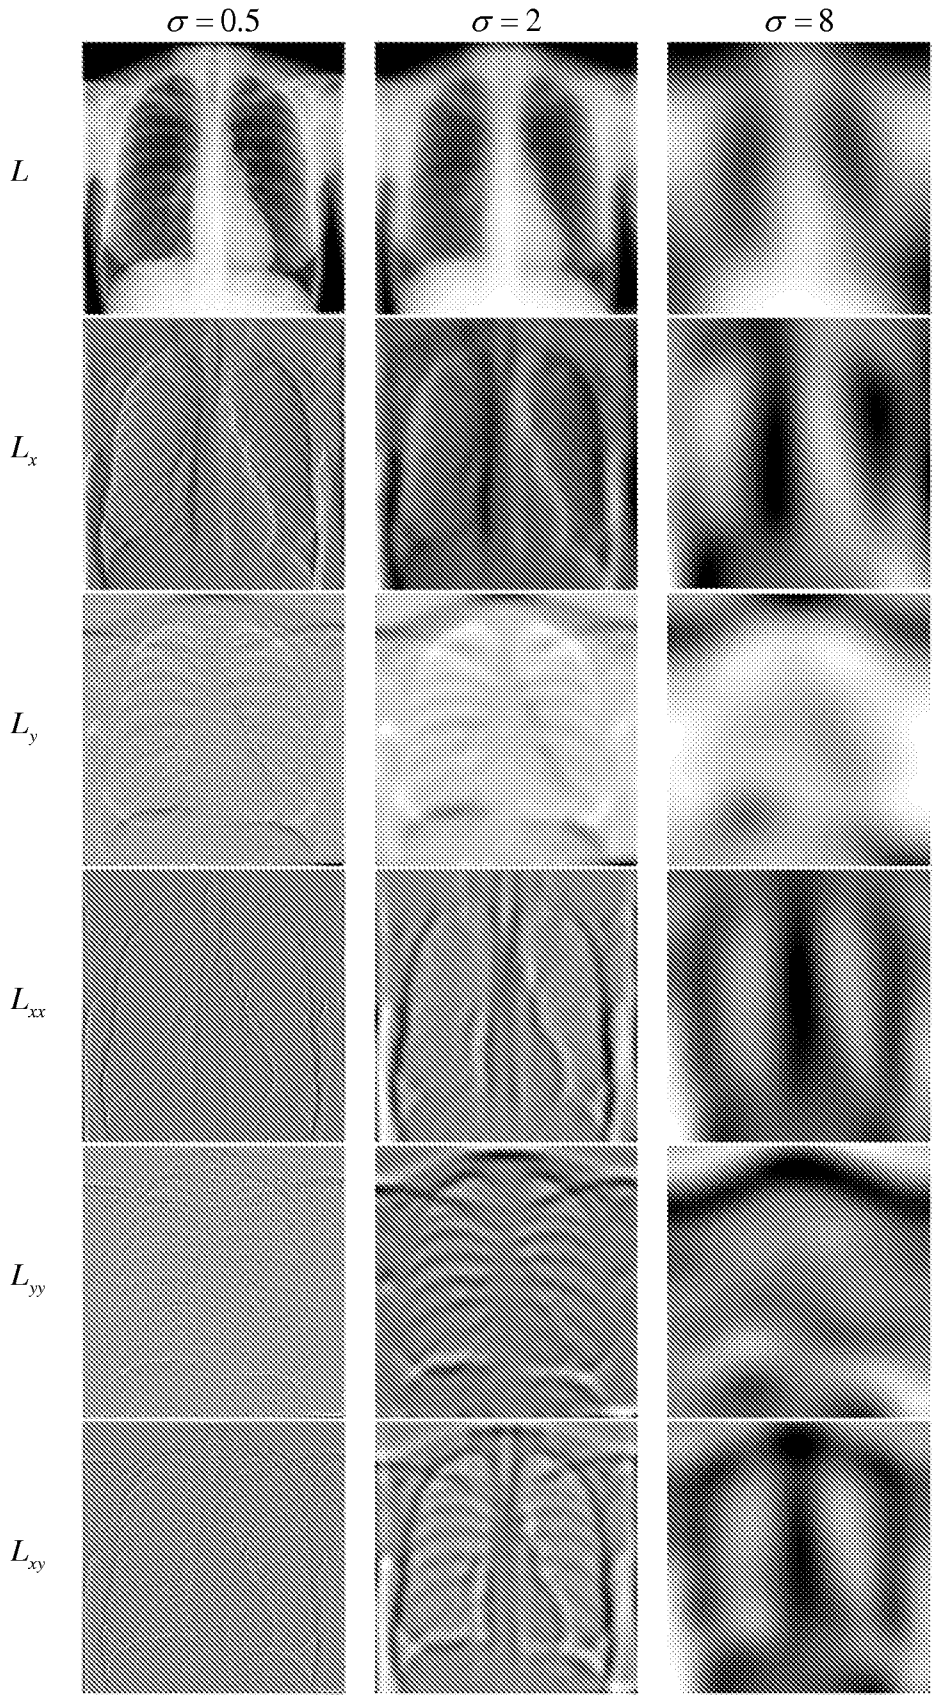


FIG. 11

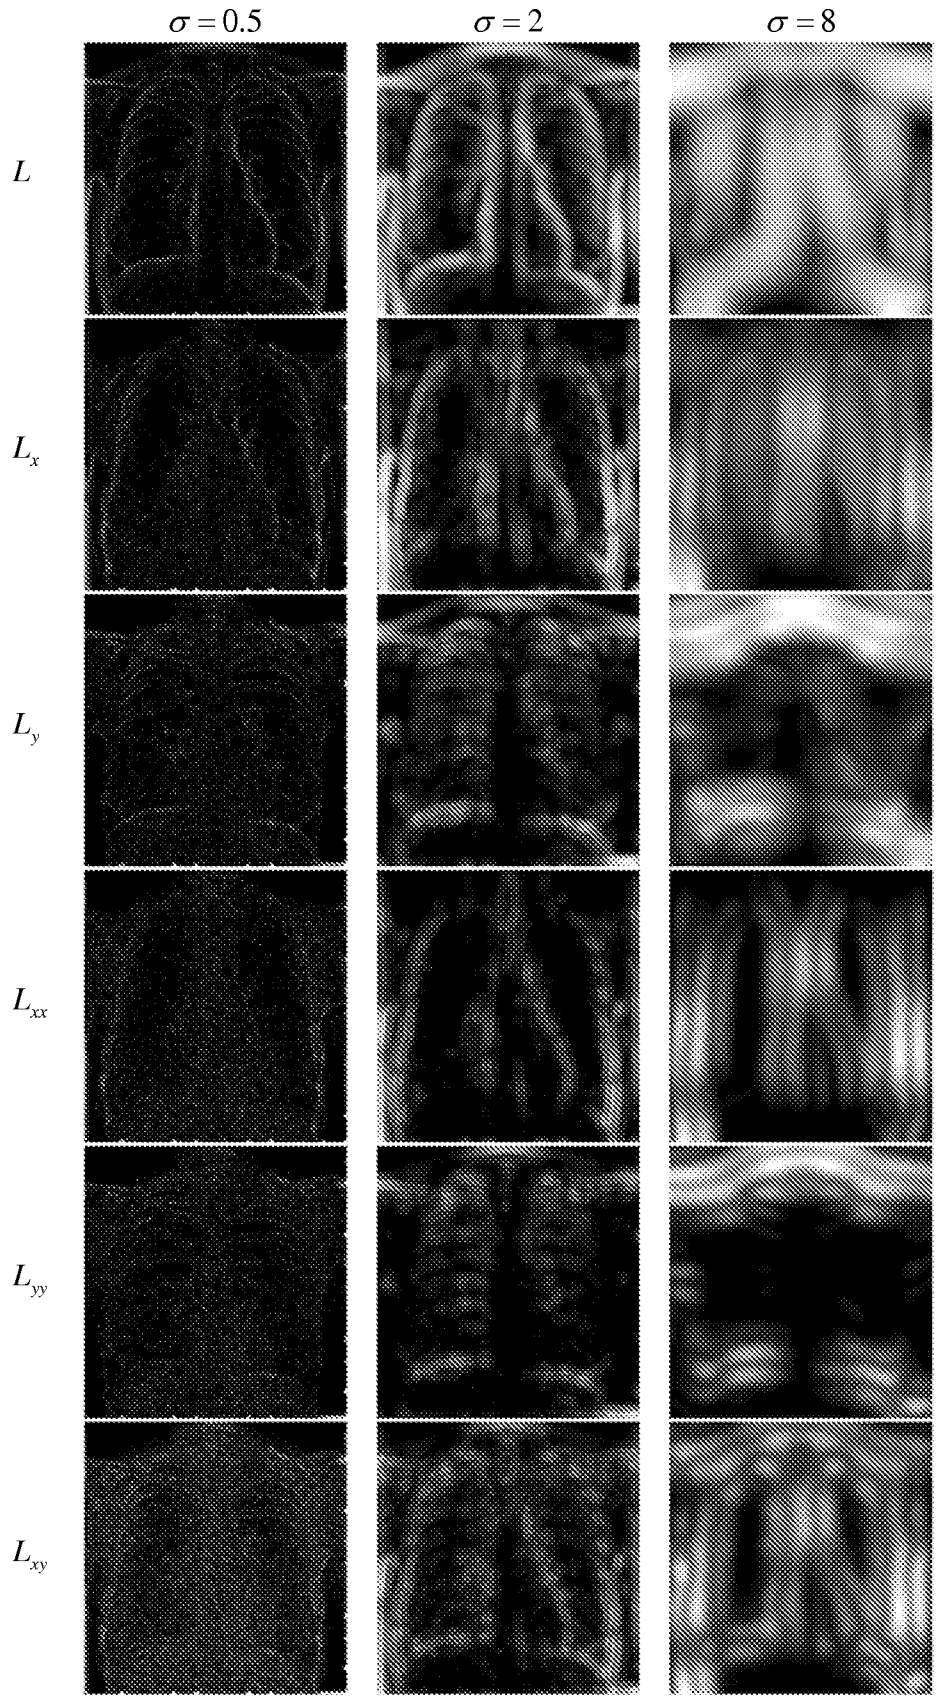


FIG.12

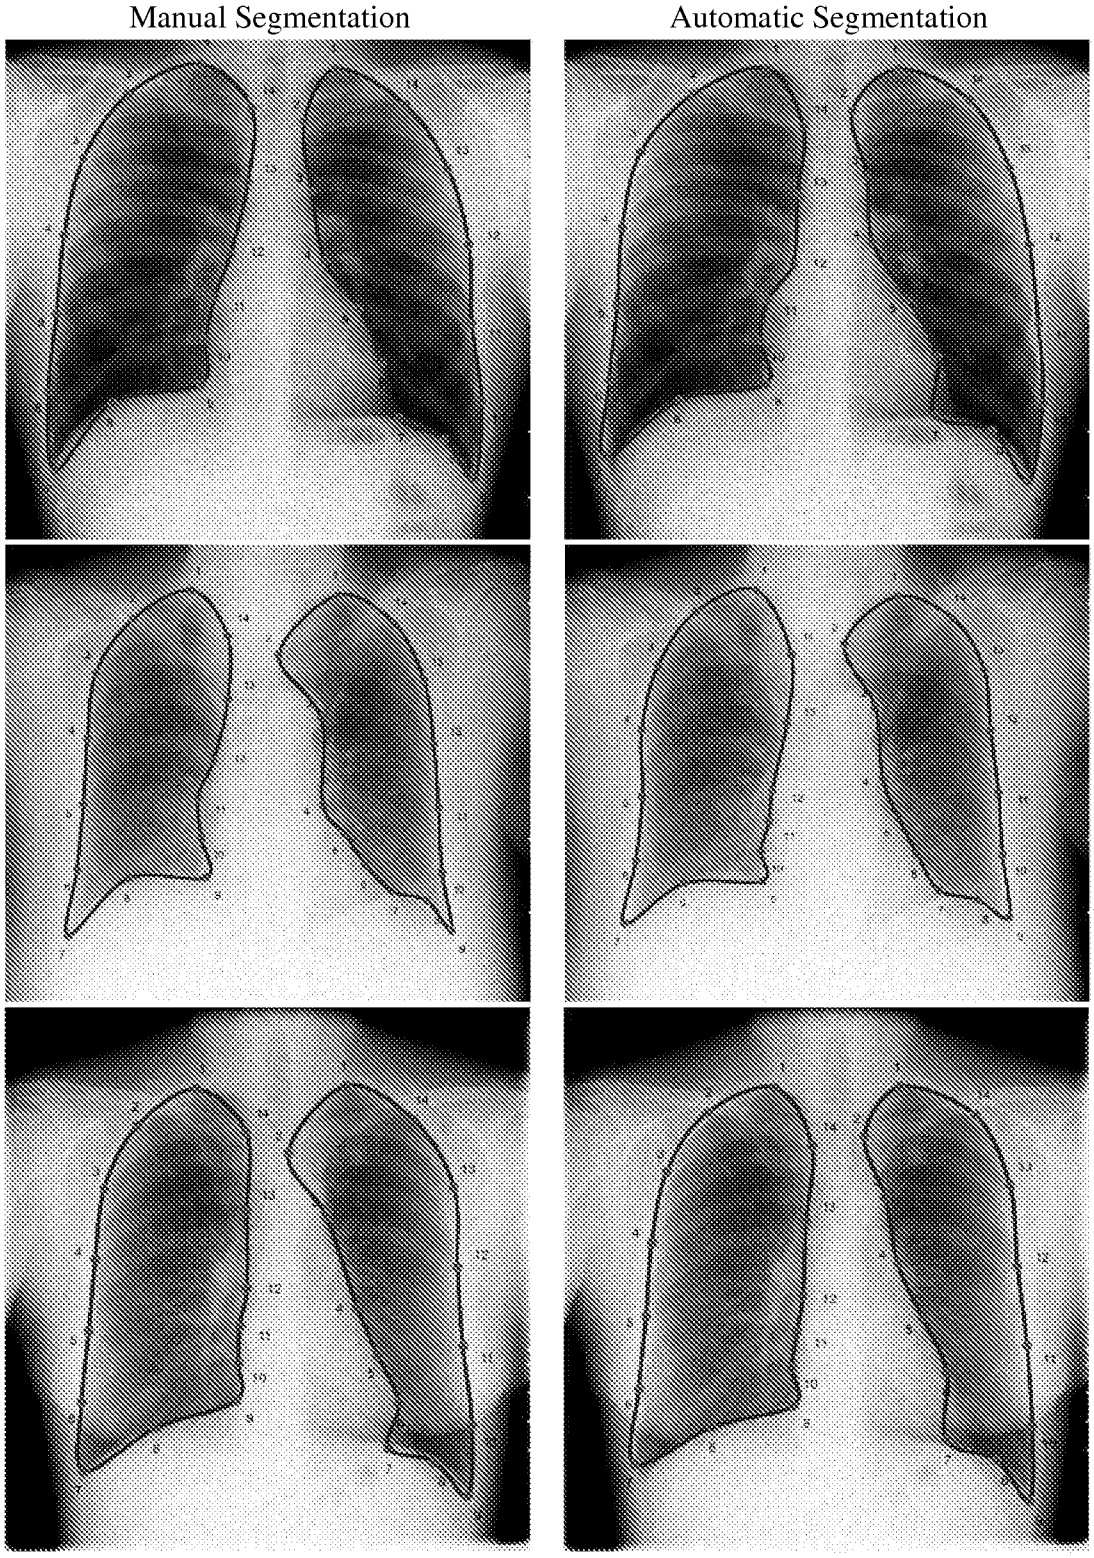
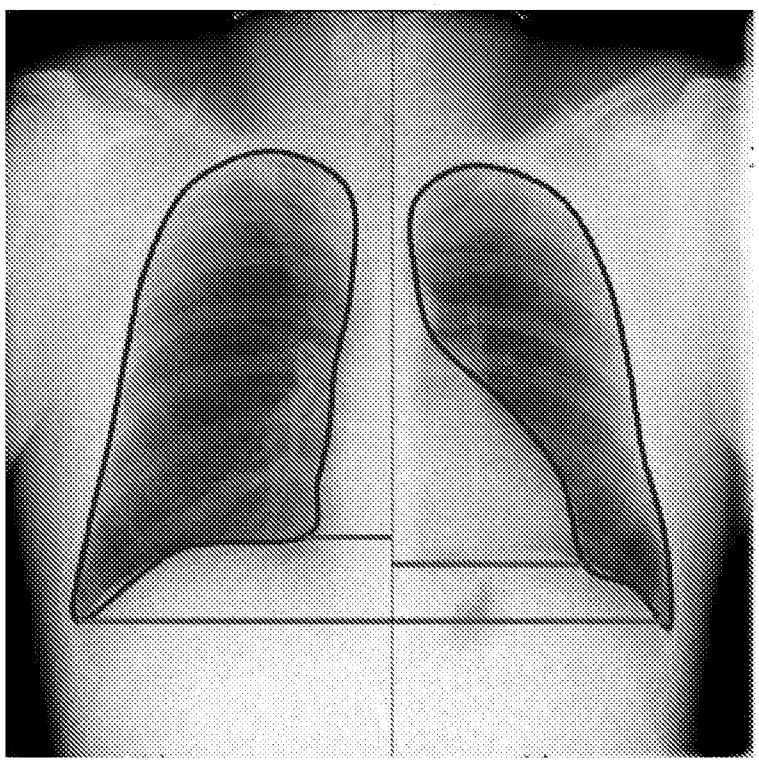
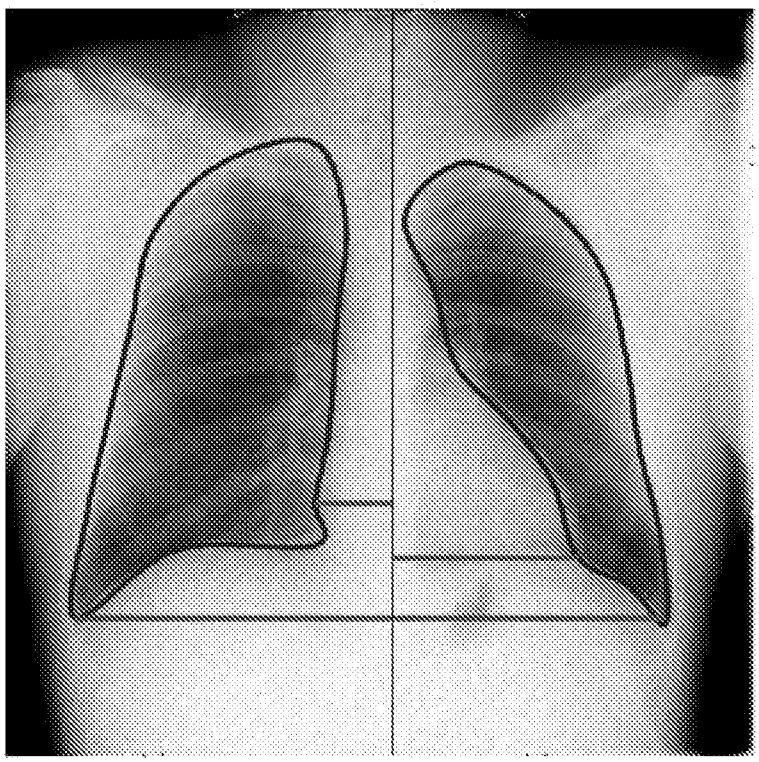


FIG. 13



(a)



(b)

FIG. 14

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2008/057188

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T5/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2007/047790 A1 (DEWAELE PIET [BE]) 1 March 2007 (2007-03-01) paragraph [0172]; figures 1-15	1-10

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
15 July 2008	23/07/2008

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Kulak, Eray

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2008/057188

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007047790 A1	01-03-2007	NONE	