

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
18 August 2011 (18.08.2011)(10) International Publication Number
WO 2011/098752 A2(51) International Patent Classification:
G06K 9/64 (2006.01)

(74) Agent: KENRICK, Mark; Marks & Clerk LLP, Sussex House, 83 - 85 Mosley Street, Manchester M2 3LG (GB).

(21) International Application Number:
PCT/GB2011/000157

(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, CL, 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, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date:
8 February 2011 (08.02.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/703,438 10 February 2010 (10.02.2010) US
1002270.5 11 February 2010 (11.02.2010) GB

(71) Applicant (for all designated States except US): IMOR-PHICS LIMITED [GB/GB]; Kilburn House, Manchester Science Park, Lloyd Street North, Manchester M15 6SE (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): VINCENT, Graham Richard [GB/GB]; 35 Dudley Road, Whalley Range, Manchester M16 8FW (GB). BOWES, Michael Antony [GB/GB]; Paddock Lodge, Paddock Lane, Whalley Bridge, Derbys SK23 7JP (GB). GUILLARD, Gwennael Alain [GB/GB]; 17 Park Drive, Stockport SK4 2HJ (GB). SCOTT, Ian Michael [GB/GB]; 39 Shaw Road, Stockport SK44AG (GB).

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

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: IMAGE ANALYSIS




Fig. 1

(57) Abstract: A computer-implemented method for identifying features in an image. The method comprises fitting a plurality of second models to the image, the plurality of second models together modelling a region of interest, wherein each part of the region of interest is modelled by at least two of the plurality of second models; and identifying the features in the image based upon the fit of the plurality of second models.

IMAGE ANALYSIS

5 The present invention relates to methods and apparatus for identifying features in an image. More particularly, but not exclusively, the present invention relates to methods and apparatus for fitting models to an image to identify features in the image.

10 Automated image analysis may be used to reduce manual workloads in determining properties of images. Image analysis is used in a variety of different fields to identify features of images. One area in which image analysis techniques are used is to process medical images where the identification of anatomical features in an image can provide useful information to medical practitioners. Often many images are required to be processed, for example in a screening program, and each image should be processed as accurately as possible.

15 Images of a particular anatomical feature obtained from different patients vary considerably. This may particularly be the case when an anatomical feature is in a diseased state. It can therefore be difficult to accurately identify anatomical features in images given such variance. An existing approach to identifying anatomical features in an image is based upon statistical modelling of an anatomical region of interest.

20 Statistical models are generated from a training set of images in which features within the anatomical region of interest have been manually identified and a statistical model indicates a mean model and an acceptable degree of variation from the mean model, as indicated by the training set. Where a small number of training images are available, a statistical model can often fail to identify particular anatomical features in previously 25 unseen images due to the failure of the statistical model to capture a true indication of acceptable variation.

30 Additionally, existing approaches for fitting statistical models to previously unseen images can be inaccurate, particularly when an anatomical feature is in a diseased state in the previously unseen image. A robust and accurate method of modelling structures in images and fitting such models to previously unseen images is therefore desirable.

It is an object of some embodiments of the present invention to obviate or mitigate at least some of the problems set out above.

According to a first aspect of the invention there is provided a computer-implemented 5 method for identifying features in an image. The method comprises fitting a plurality of second models to the image, the plurality of second models together modelling a region of interest, wherein each part of the region of interest is modelled by at least two of the plurality of second models; and identifying the features in the image based upon the fit of the plurality of second models.

10

Fitting a plurality of second models to an image allows localised regions to be individually modelled. The local nature of the second models can allow a more accurate fit to the image. By including parts of the region of interest to be modelled in more than 15 one of the second models, if the fit of any one of the second models to a particular region is unacceptable, another of the second models that models the particular region can be used.

The method may further comprise initially fitting a first model of the object to the 20 image and fitting the plurality of second models may be based upon the fitting of the first model.

Fitting the plurality of second models to the image may be initialised based upon the fitting of the first model. Initialising the fitting of models based upon a fit of a first model or models in this way allows information determined from fitting a first model to 25 guide a fitting of subsequent models. Alternatively or additionally fitting the plurality of second models may be limited by the fitting of the first model.

Identifying features in the image based upon the fit of the plurality of second models may comprise combining data based upon the fit of the plurality of second models.

30

Combining data based upon the fit of the plurality of second models may comprise, for each part of the region of interest included in at least two of the plurality of models, applying an associated weight to data derived from each of the plurality of second

models in which the part is included. For a particular second model, the associated weight may be based upon a distance of the part of the region from a point defined with respect to the particular second model. The point defined with respect to the particular second model may be defined with respect to a centre of the particular second model, 5 for example the centre of a region modelled by a particular second model.

A particular feature may be modelled by two of the second models and the combining may identify the particular feature by combining data associated with the particular feature generated by fitting each of the two of the second models. The combining may 10 be based upon stored data indicating a plurality of relationships, each relationship being between at least part of a second model and at least part of at least one other second model. For example, each point in a particular second model may have a relationship with any point in any other second model that represents a corresponding point in a region of interest. The stored data indicating a plurality of relationships may indicate a 15 relationship between at least part of a second model and part of the first model.

The combining may comprise, for each part of the region of interest included in at least two of the plurality of second models, determining a weighted average fit for the part of the region of interest based upon the fit of each of the at least two of the plurality of 20 second models to the part of the region of interest.

Identifying the features in the image based upon the fit of the plurality of second models may further comprise determining an estimate of location of the features in the image based upon the fit of the plurality of second models. Determining an estimate of 25 location of features in the image may comprise obtaining an initial estimate of location of the features in the image and updating the initial estimate based upon the fit of the plurality of second models. In this way, an initial estimate of location of the features in an image may be refined based upon subsequent fits of models

30 Updating the initial estimate may comprise updating the initial estimate based upon the fit of one of the second models if the fit of the one of the second models satisfies a predetermined criterion. The predetermined criterion may be a relationship between the initial estimate of location of the features and the fit of the one of the second models.

The predetermined criterion may comprise a geometric criterion. The predetermined criterion may be determined for the one of the second models during a training phase, based upon a set of training images. The predetermined criterion may be based upon a plurality of data items, each data item representing a fit of the one of the second models 5 to a respective one of the training images. Each of the plurality of second models may have an associated predetermined criterion.

The method may further comprise fitting a plurality of third models to the image, the plurality of third models together modelling at least part of the region of interest, the 10 fitting the plurality of third models being based upon the estimate of location of the features, and identifying the features in the image based upon the fit of the plurality of third models. The estimate of location of the features may be used to initialise the third models in the image. The estimate of location of the features may be updated based upon the fit of the plurality of third models. Features in the image may be identified 15 based upon the updated estimate of location of features. Further pluralities of models may be fitted to the image based upon the estimate of location of the features, with the estimate of location of the features updated after fitting each plurality of models.

A given part of the region of interest may be modelled by a first number of the second 20 models and a second different number of the third models. The second number may be greater than the first number. Models in the third plurality of models may generally model smaller regions than models in the second plurality of models. In this way increasingly localised regions of the image may be modelled and a more accurate fit to the image may be achieved.

25 The method may further comprise generating the plurality of second models. A training set of images may be received, each image in the training set of images representing the region of interest, and generating the plurality of second models of the object may be based upon the training set of images. Each model in the plurality of second models may be 30 generated based upon a subset of images in the training set of images, the subset being selected by randomly sampling with replacement from the training set of images.

Generating the plurality of second models may comprise generating a plurality of models, each of the plurality of models modelling at least part of the region of interest; processing the plurality of models with respect to a further training set of images; and selecting a subset of the plurality of models based upon the processing such that each 5 part of the region of interest is represented by a plurality of models in the subset of models.

Generating the plurality of second models may further comprise generating a first plurality of second models such that each part of the region of interest is modelled by a 10 plurality of the first plurality of second models; generating a second plurality of second models such that each part of the region of interest is modelled by a plurality of the second plurality of second models; processing the first plurality of second models and the second plurality of second models with respect to the further training set of images; and selecting one of the first plurality and the second plurality of second models of the 15 region of interest based upon the processing. Processing the first plurality of second models and the second plurality of second models with respect to a training set of images may comprise fitting the first plurality of second models to each image in the further training set of images; and fitting the second plurality of second models to each image in the further training set of images; and selecting one of the plurality of second 20 models of the region of interest based upon the processing may be based upon the fittings.

Generating the or each plurality of second models may comprise receiving data indicating a plurality of points of interest within the region of interest; and generating the or each plurality of second models based upon the received plurality of points of 25 interest. Generating the or each plurality of second models based upon the received data may comprise, for each of the received points of interest, selecting a part of the region of interest of predetermined size to provide a basis for one of the plurality of models. Each part of the region of interest may represent a feature of interest.

30

Generating the plurality of second models of the object based upon the training set of images may comprise, for each of the plurality of points, processing a corresponding part of the region of interest in each image in the training set of images.

According to a second aspect of the invention there is provided a computer-implemented method for identifying features in an image. The method comprises fitting a plurality of second models to the image, the plurality of second models together modelling a region of interest; fitting a plurality of third models to the image, the plurality of third models together modelling at least part of the region of interest, the fitting of the plurality of third models being based upon the fit of the plurality of second models; and identifying the features in the image based upon the fit of the plurality of third models.

10

In this way, a fit to the image in a given region may be modified based upon fits of different sets of models. Such a method provides a way of refining identification of an object in an image. Further pluralities of models may be fitted to the image based upon previous pluralities of models to further refine the identification.

15

A given part of the region of interest may be modelled by a first number of second models and a second different number of third models. The second number may be greater than said first number. Some or all of the third number of models may model a smaller region than some or all of the second number of models and so each of the third 20 plurality of models may model a more localised part of the region of interest. Improved modelling of the region of interest may thereby be achieved.

25

The method may further comprise initially fitting a first model of the region of interest to the image and fitting the plurality of second models may be based upon the fitting of the first model. Fitting the plurality of second models to the image may be initialised based upon the fitting of the first model and fitting the plurality of third models to the image may be initialised based upon the fitting of the first model. As indicated above, initialising models based upon a fit of previous models in this way allows information determined from fitting the previous models to guide the subsequent fitting.

30

Each region of the object may be represented by a plurality of the plurality of second models. Advantages described for the first aspect of the invention may therefore also be achieved in the second aspect of the invention.

According to a third aspect of the invention there is provided a computer-implemented method for identifying features in an image. The method comprises determining an estimate of location of the features in the image; fitting a second model to part of the image; and updating the estimate based upon the second model if the fit of the second model satisfies a predetermined criterion.

In general terms, it is often possible to provide a reasonably good estimate of the location of features in an image. Fitting a second model to part of the image allows the part of the image to be more accurately modelled. However if the second model indicates that the initial estimate should be updated by a large amount then it may be assumed that the fitting of the second model has failed and the estimate may therefore not be updated. As such, the predetermined criterion may be based upon a difference between the estimate and the fit of the second model.

The predetermined criterion may be a relationship between the estimate of location of the features and the fit of the second model. The predetermined criterion may comprise a geometric criterion, for example a maximum distance between the points modelled by the second model after fitting and corresponding points in the estimate of location of the features or a maximum amount of rotation of part of the image modelled by the second model and a corresponding part of the estimate of location, or a combination of both. The predetermined criterion may be determined for the second model during a training phase, based upon a set of training images. The predetermined criterion may be based on a plurality of data items, each data item representing a fit of the second model to a respective one of said training images. That is, the second model may be fitted to a plurality of training images and the fit of the second model to each training image may be classified as a “good” or “bad” fit. The classified fits may be used to train a classifier for determining whether a particular fit of a second model is an acceptable fit.

The method may further comprise fitting a first model to the image and the estimate of location of the features in the image may be based upon the fit of the first model.

Fitting the second model to part of the image may be based upon the estimate of location of the features. Additionally, fitting the second model to part of the image may be initialised based upon the estimate of location of the features.

5 Fitting a second model to part of the image may comprise fitting a plurality of second models to the image, and each of the plurality of second models may have an associated predetermined criterion. That is, a criterion for a particular second model may allow the particular second model to vary a relatively large amount from a starting position compared to a different second model.

10

According to a fourth aspect of the invention there is provided a computer-implemented method of generating a plurality of second models that together model a region of interest. The method comprises generating the plurality of second models such that each part of the region of interest is modelled by a plurality of the plurality of second models.

15

As indicated above, by including regions of a region of interest to be modelled in more than one of the second models, if the fit of any one of the second models to a particular region is unacceptable, another of the second models that models the particular region can be used when the second models are used to model the region of interest.

20

The method may further comprise receiving a training set of images, each image in the training set of images representing the region of interest and generating the plurality of second models of the object may be based upon the training set of images. Each model in the plurality of second models may be generated based upon a subset of images in the training set of images, the subset being selected by randomly sampling with replacement from the training set of images.

25

Generating the plurality of second models may comprise: generating a plurality of models, each of said plurality of models modelling at least part of the region of interest; processing the plurality of models with respect to a further training set of images and selecting a subset of the plurality of models based upon said processing such that each part of the region of interest is represented by a plurality of models in the subset of models.

30

Given that models are determined in an initial stage, and are generally only determined once for a particular region of interest, a relatively large amount of time can be spent in optimising the selected models. As such, many different models can be generated in any 5 convenient way and the efficacy of each generated model in fitting to the region of interest can be tested using a further training set of images. Models found to optimally fit to a region of interest may be retained while other models are rejected.

Generating the plurality of second models may comprise: generating a first plurality of 10 second models such that each part of the region of interest is modelled by a plurality of the first plurality of second models; generating a second plurality of second models such that each part of the region of interest is modelled by a plurality of the second plurality of second models; processing the first plurality of second models and the second plurality of second models with respect to the further training set of images; and 15 selecting one of the first plurality of second models and second plurality of second models of the region of interest based upon the processing.

Processing the first plurality of second models and the second plurality of second models with respect to a training set of images may comprise: fitting the first plurality 20 of second models to the object; and fitting the second plurality of second models to the object; and selecting one of said plurality of second models of the object based upon the processing may be based upon the fittings.

Generating the or each plurality of second models may comprise receiving data 25 indicating a plurality of points of interest within the region of interest and generating the or each plurality of second models based upon the received plurality of points of interest. The points of interest may be selected in any convenient way. For example, each point of interest may be indicative of a feature of interest such as, for example, features of the region of interest and/or regions known to be of greatest curvature in the 30 region of interest. In this way, particular features which it is particularly desirable to accurately model can be selected so as to fall at the centre of a model, and therefore be more accurately modelled.

Generating the or each plurality of second models based upon the received data may comprise, for each of the received points of interest, selecting a part of the region of interest of a predetermined size to provide a basis for one of the plurality of models. Each part of the region of a predetermined size may be centred on one of the received 5 points of interest.

Generating the plurality of second models of the object based upon the training set of images may comprise, for each of the plurality of points, processing a corresponding part of the region of interest in each image in the training set of images.

10

In each of the aspects the models may be any suitable description of shape and/or image characteristics which can be fitted to an image. For example, the models may be rigid models that may be fitted to an image by changing the position and/or rotation of the models to determine a best fit to the image. Alternatively the models may be deformable 15 models. Deformable models are models in which shape of a region may be changed within predetermined constraints in order to effectively fit the models to the image. Examples of deformable models are statistical models such as shape or appearance models, deformable templates and deformable atlases. Where the models are appearance models they may represent a volume or a surface.

20

The or each image may be an image selected from the group consisting of: a computed topography image, a magnetic resonance image, a radiograph image, an x-ray image and an ultrasound image. The fitting may comprise applying a fitting technique selected from the group consisting of: rigid registration, non-rigid registration, active shape 25 modelling and active appearance modelling. The plurality of second models may represent an arbitrary decomposition of the object. The region of interest may be a part of a body.

The term “features” is used herein to indicate any part of an image which is to be 30 identified. For example a feature may be an image region, an image boundary, an image surface or a point in an image. In the case where the region of interest is a part of a body the features may be anatomical features. Examples of anatomical image regions include a region indicative of bone such as a femur, a region indicative of an organ or part of an

organ (for example a region indicative of the frontal lobe), regions indicative of abnormalities (for example regions indicative of cysts or regions indicative of pathological bone growth). Examples of anatomical boundaries and surfaces include bone/cartilage boundaries and/or surfaces and ventricle/cerebrospinal fluid boundaries and/or surfaces. An example of an anatomical point in an image is an iliac crest.

Aspects of the invention can be implemented in any convenient form. For example computer programs may be provided to carry out the methods described herein. Such computer programs may be carried on appropriate computer readable media which term includes appropriate tangible storage devices (e.g. discs). Aspects of the invention can also be implemented by way of appropriately programmed computers.

Features described in the context of one aspect can be used in the context of another aspect.

15

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a schematic illustration of a system for analysis of images according to an embodiment of the present invention;

Figure 1A is a schematic illustration showing a computer of the system of Figure 1 in further detail;

25

Figure 2 is an image suitable for processing using the system of Figure 1;

Figure 3 is a flowchart showing processing carried out to analyse an image in the system of Figure 1;

30

Figure 4 shows the image of Figure 2 divided into parts, each of which is modelled by a respective model;

Figure 5 is a flowchart showing part of the processing of Figure 3 in further detail;

Figure 6 is a flowchart showing part of the processing of Figure 5 in further detail; and

5 Figure 7 is a graph illustrating a relationship between a weight applied to a point and a distance of the point from the centre of a model.

Referring now to Figure 1, a computer 1 is arranged to receive an image 2 of an object or objects and to receive a model 3. The model 3 is fitted to the image 2 to generate a parameterisation of the model 4. The model 3 is a statistical model such as an active 10 appearance model or active shape model, which has been created based upon a training set of images of an object of the same class as the object that is the subject of image 2. The model 3 is a statistical model of the variation within the training images and is generated by processing the training set of images to generate a mean model and a range of variation from the mean model of the training set. The parameterisation 4 generated 15 by computer 1 can be used to identify features of the image.

Figure 1A shows the computer 1 in further detail. It can be seen that the computer comprises a CPU 1a which is configured to read and execute instructions stored in a volatile memory 1b which takes the form of a random access memory. The volatile 20 memory 1b stores instructions for execution by the CPU 1a and data used by those instructions. For example, in use, the image 2 and model 3 may be stored in the volatile memory 1b.

The Computer 1 further comprises non-volatile storage in the form of a hard disc drive 25 1c. The image 2 and model 3 may be stored on the hard disc drive 1c. The computer 1 further comprises an I/O interface 1d to which is connected peripheral devices used in connection with the computer 1. More particularly, a display 1e is configured so as to display output from the computer 1. The display 1e may, for example, display a representation of the image 2. Additionally, the display 1e may provide data representing the parameterisation 4 to a user and may display an input image 2 overlaid 30 with the parameterisation 4. Input devices are also connected to the I/O interface 1d. Such input devices include a keyboard 1e and a mouse 1f which allow user interaction with the computer 1. A network interface 1g allows the computer 1 to be connected to

an appropriate computer network so as to receive and transmit data from and to other computing devices. The CPU 1a, volatile memory 1b, hard disc drive 1c, I/O interface 1d, and network interface 1g, are connected together by a bus 1h.

5 Referring now to Figure 2 an image suitable for processing by the system of Figure 1 is shown. Figure 2 is a slice through a Magnetic Resonance (MR) volume image of the knee. The system of Figure 1 is able to process different kinds of images in the same way. For example, images derived from annotated data such as three-dimensional surfaces or images of anatomical volumes can be processed in the manner described
10 below.

The processing described below makes reference to a model M_1 , built based upon a training set of images. The model M_1 is a statistical model built based upon a training set of images and is a model of features of interest in an image I , as set out above. In
15 order to aid the understanding of the description of the invention below, generation of the model M_1 and fitting of model M_1 to an image I will first be described. It will be appreciated however that the generation and fitting of the model M_1 can be carried out in any convenient way. For example, the model M_1 can be generated using a parameterisation technique described in United States Patent Number 7584080, the
20 contents of which are herein incorporated by reference.

The model M_1 is generated from a training set of images by first manually marking corresponding predetermined points of interest in each of the images in the training set of images. The set of points of interest are selected based upon features of the object to be modelled. For example, if the object to be modelled is a knee joint, points of interest may be selected based upon anatomical features of the knee joint. The model M_1 may be, for example, a shape model or an appearance model, which are both described below. It will be appreciated, however, that other models may be used such as for example, deformable templates or deformable atlases which take the form of a typical or
25 average example which is allowed to deform within a set of constraints. A review of deformable models is given in Tim McInerney and Demetri Terzopoulos "Deformable Models in Medical Image Analysis: A Survey", Medical Image Analysis, 1(2):91–108,
30 1996 the contents of which is herein incorporated by reference.

5 In the case where the model M_1 is a shape model, the set of points of interest associated with each training image, selected in the manner described above, is represented by a shape vector comprising a concatenation of the coordinates of the selected points. The shape vector \mathbf{x} for a particular training image takes the form:

$$\mathbf{x} = (x_1, y_1, x_2, y_2, \dots, x_n, y_n) \quad (1)$$

where n is the number of points of interest in each image.

10

A mean shape vector $\bar{\mathbf{x}}$ can be calculated according to equation (2):

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i \quad (2)$$

Where N is the number of training images; and

15

\mathbf{x}_i is a vector representing a training image i

A plurality of shape vectors, one for each training image can be used to generate a matrix \mathbf{X} , in which each row represents a shape vector.

20

Each of the shapes represented by the points selected in a respective training image is translated, rotated and scaled so that their centroids coincide and the sum of squared differences between the points in each image and those on the mean shape is minimised.

25

A covariance matrix \mathbf{S} can be created from the aligned shapes by first calculating the mean shape vector using equation (2). For each image in the training set, the deviation from the mean shape can be calculated from equation (3):

$$d\mathbf{x}_i = \mathbf{x}_i - \bar{\mathbf{x}} \quad (3)$$

30

The covariance matrix \mathbf{S} is then calculated using equation (4):

$$S = \frac{1}{N} \sum_{i=1}^N d\mathbf{x}_i d\mathbf{x}_i^T \quad (4)$$

The eigenvectors of the matrix S can be determined and a matrix P of the t most significant eigenvectors can be created.

5

In this way, any shape can be represented according to the shape model given by equation (5):

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}_s \mathbf{b}_s \quad (5)$$

10

where \mathbf{x} is a vector representing the shape;

$\bar{\mathbf{x}}$ is a mean shape vector generated from shape vectors of images in the training set (according to equation (2));

\mathbf{P}_s is the matrix of eigenvectors described above; and

15

\mathbf{b}_s is a t -element vector of weights to be applied to eigenvectors of the matrix \mathbf{P}_s to generate the vector \mathbf{x} for the shape.

20

The components of \mathbf{b}_s (b_i , $i = 1 \dots t$) are referred to as shape parameters. The eigenvectors \mathbf{P}_i are sorted according to their eigenvalues λ_i , which describe the observed variance over the training shapes associated with each eigenvector. That is to say the λ_i represent the proportions of variance from the mean shape that are described by each of the eigenvectors. By selecting a value of t ($\leq N$), the model can be made to represent some chosen proportion of the total variance observed in the training shapes.

25

Each eigenvector included in the matrix \mathbf{P}_s represents a mode of variation of the modelled shape. Varying values of elements of the vector \mathbf{b}_s allows generation of shapes within the range observed within images of the training images. That is, variation of one element b_1 of the vector \mathbf{b}_s will affect one characteristic of the modelled shape, while variation of a second element b_2 will affect a different characteristic of the modelled shape.

30

An appearance model is an extension of the shape model described above. To build a statistical model of the grey-level appearance after the points selected in a respective training image are translated, rotated and scaled so that their centroids coincide and the sum of squared differences between the points in each image and those on the mean shape is minimised, as described above, grey-level information g_{im} is sampled from the image over the region covered by the mean shape \bar{x} , after normalising the image to minimise the effect of global lighting variation. Applying the method set out above to determine the shape model (5) to the grey-level data, a grey-level model is obtained, as shown in equation (6):

10

$$\mathbf{g} = \bar{\mathbf{g}} + \mathbf{P}_g \mathbf{b}_g \quad (6)$$

where:

15 $\bar{\mathbf{g}}$ is the mean normalised grey-level vector;
 \mathbf{P}_g is a matrix of eigenvectors; and
 \mathbf{b}_g is a set of grey-level parameters.

Varying values of elements of the vectors \mathbf{b}_s (in equation (5)) and \mathbf{b}_g (in equation (6)) allows generation of shapes and grey-level images within the range observed within 20 images of the training images.

In order that the relationship between shape and grey-level variation in the training images is captured, a vector \mathbf{b} is generated for each training image according to equation (7):

25

$$\mathbf{b} = \begin{pmatrix} \mathbf{W}_s \mathbf{b}_s \\ \mathbf{b}_g \end{pmatrix} = \begin{pmatrix} \mathbf{W}_s \mathbf{P}_s^T (\mathbf{x} - \bar{\mathbf{x}}) \\ \mathbf{P}_g^T (\mathbf{g} - \bar{\mathbf{g}}) \end{pmatrix} \quad (7)$$

where:

30 \mathbf{W}_s is a diagonal matrix of weights associated with each shape parameter in the vector \mathbf{b}_s .

The matrix \mathbf{W}_s adjusts values of \mathbf{b}_s to allow for the difference in units between the shape and grey level models.

A further model shown in equation (8) is generated from the set of concatenated vectors \mathbf{b} , using the method described above to generate models (5) and (6).

5

$$\mathbf{b} = \mathbf{Q}\mathbf{c} \quad (8)$$

where:

\mathbf{Q} is a matrix of eigenvectors; and
 \mathbf{c} is a vector of appearance parameters.

10

Since the model of equation (8) is linear, shape and grey-levels may be expressed directly as functions of \mathbf{c} according to equation (9).

15

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}_s \mathbf{W}_s \mathbf{Q}_s \mathbf{c} \quad , \quad \mathbf{g} = \bar{\mathbf{g}} + \mathbf{P}_g \mathbf{Q}_g \mathbf{c} \quad (9)$$

where the matrix of eigenvectors $\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_s \\ \mathbf{Q}_g \end{pmatrix}$.

An example image can be synthesised for a given \mathbf{c} by generating the shape-free grey-level image from the vector \mathbf{g} and warping it using the control points described by \mathbf{x} .

20

In general terms, the aim of model fitting is to determine the set of parameters \mathbf{c} which best fit the model to an unseen image. One way to do this is by using active appearance models which are described below.

25

An active appearance model uses the difference between a reconstructed image generated by the model and the underlying target image \mathbf{I} , to iteratively drive the model parameters towards better values. In a prior learning stage, known displacements to the vector of appearance parameters, $\delta\mathbf{c}$, are applied to known model instances and the resulting difference between model and image, $\delta\mathbf{I} = \mathbf{I}_i - \mathbf{I}_m$, is measured, where \mathbf{I}_i is a vector of grey-level values in the image \mathbf{I} and \mathbf{I}_m is a vector of grey-level values for the current model parameters \mathbf{c} . Multivariate linear regression is applied to a large set of

30

such training displacements and an approximate linear relationship, $\delta c = A\delta I$, is established.

To identify the best match between the model and image I the magnitude of the difference vector $\Delta = |\delta I|^2$ is minimised using the linear relationship learnt during the learning stage to predict an adjustment to the model parameters c which improves model fit. The model parameters c are repeatedly adjusted until no improvement is made to Δ . Active appearance models are described in detail in T.F. Cootes, C.J. Edwards, C.J. Taylor "Active Appearance Models", IEEE Transactions on Pattern Analysis and Machine Intelligence, 2001, Volume 23, Issue 6, which is herein incorporated by reference.

Referring now to Figure 3, processing to fit a model to an image is shown. At step S1 an image I is received and at step S2 a model M_1 is fitted to the image I. The model M_1 is a statistical model built based upon a training set of images and is a model of features of interest in the image I such as the models described above.

At step S3 a counter k is set to a value 2 and at step S4 each model m of a set of models M_k is initialised in the image I based upon the fit of the model M_1 . The set of models M_2 is a decomposition of an anatomical region modelled by the model M_1 . That is, each point modelled by each model m corresponds to a point modelled by the model M_1 , and collectively the set of models M_2 model all of the features of interest modelled by the model M_1 in the image I. For all values of k , models of the set of models M_2 overlap such that each point of the object in the image I that is modelled by model M_1 is modelled by a plurality of the models m in the set of models M_2 . Decomposition of the anatomical region modelled by the model M_1 to generate a set of models M_k is described in further detail below with reference to Figure 4.

Initialisation of each model m is performed by fitting points in each model to the image I based upon the fit of corresponding points of the model M_1 , determined at step S2. Corresponding points are determined using a point correspondence indicating a relationship between points. The point correspondence relates each point modelled by the model M_1 to points in models m (for all M_k) that correspond to the respective point.

In this way, the relationship between points in each model in a set of models M_k is stored, as well as the relationship between points in models m_k in a set of models M_k and models m_{k+1} in a set of models M_{k+1} .

5 At step S5 the fit of each model m to the image I is adjusted based upon properties of the image I and properties of the model m in the same way as the fitting of model M_1 described above. Each model m is adjusted independently of any other model m in the set of models M_k and each model m generates an estimate of location of features in the image I that are modelled by m . At step S6 the fits of the models in the set M_k to the
10 image I after the adjustment of step S5 are combined to generate a combined estimate which provides an estimate of the location of each of the features of interest in the image. The combined estimate may take the form of any suitable data. For example the combined estimate may be a geometric indication of the location of each of the features of interest in the image I , or may be a description of the image I indicating the features of interest in the image I , or a combination of both. Combining the fits of the models M_k to generate a combined estimate is described in further detail below with reference to
15 Figure 5.

At step S7 a check is performed to determine whether there are more sets of models M_k to be fitted to the image. If it is determined that there are more sets of models then at
20 step S8 the counter k is incremented and processing passes to step S4 where the further set of models M_{k+1} is initialised based upon the combined estimate, determined at step S6. Otherwise at step S9 an estimate of location of features of interest is output based upon the combined estimate generated from the set of models M_k . Each set M_{k+1} generally contains a greater number of models than the previous set M_k and, in general, models in the set M_{k+1} will model a smaller number of points than models in the set M_k . As such, as k increases models in the set M_k will generally model a more localised region of the object and smaller regions of the object will generally be modelled with
25 greater accuracy.

30

Initialisation of model fitting based upon the combined estimate minimises the risk of a model incorrectly identifying part of the image I as a feature of interest when such a part of the image I is, in fact, not a feature of interest. That is, where a model models a

relatively small part of an image, there is a greater chance such a model will fit an incorrect part of an image well. As such, smaller models need better initialisation to avoid false fitting. Use of a hierarchical scheme allows a plurality of models of differing sizes to be applied to an image, such that fitting of relatively large models initializes the 5 fitting of relatively smaller models. Models M_{i+1} at each level of the hierarchy update an estimate generated by a level M_i of the hierarchy.

It has been described above that an anatomical region modelled by a model M_1 is 10 decomposed, and the decomposition is processed to generate a set of models M_k . Figure 4 illustrates a decomposition of an anatomical region modelled by a model M_1 based upon the image of Figure 2. Figure 4 shows three regions 10, 11, 12. Each of regions 10, 11, 12 corresponds to part of the image of Figure 2. Each of regions 10, 11, 12 corresponds to a subset of the points of interest modelled by model M_1 and indicates 15 regions of the image of Figure 2 that correspond to one of the models in the set of models M_k . It can be seen that regions 10, 11, 12 overlap so that some points fall within two or three of regions 10, 11, 12. For example image feature 13 of region 10 corresponds to features 14 and 15 of regions 11 and 12 respectively, such that any points corresponding to features 13, 14 and 15 that are modelled by model M_1 will be modelled by each of the models corresponding to image regions 10, 11 and 12. It has 20 been found that modelling each image point with between 5 and 15 models in the set of models M_k is effective, although each image point may be modelled in other numbers of models greater than or equal to 2, in a set M_k .

In general terms, each set of models M_k is determined by selecting a plurality of points 25 of interest and identifying a region based upon each of the selected points. The set of points selected to form the basis for a model m of the set of models M_k is processed with respect to a training set of images to generate the model m in the manner described above. It will be appreciated that the training set of images used can be the same training set as is used to generate the model M_1 , or could be a further training set of 30 images which have been parameterised with corresponding points to the points used in the generation of the model M_1 . In some embodiments a subset of the set of training images may be used to generate each model with each subset selected by randomly sampling with replacement from the set of training images. That is, each model may be

generated based upon a randomly selected subset of the training images, wherein the subset selected for generation of a particular model does not affect subsets selected for generation of other models.

5 The particular points of interest modelled by any model m in a set of models M_k may be selected in any convenient way. For example a segmenter may be used to divide the set of points of interest modelled by an initial model into overlapping subsets. The subsets may be determined by selecting points at random and determining an area of a predetermined size surrounding each selected point in each training image to provide
10 basis for a model. That is, any points falling within the determined area may then be used as a basis for a model of the set of models M_k . Points may be chosen in the same way until a desired degree of overlap has been achieved for each of the points to be modelled.

15 Alternatively, a set of points may be selected that are evenly distributed across an initial model, such that the distance between the selected points is approximately equal. An area surrounding each of the selected points may then be selected and any points of the initial model falling within the area associated with a selected point may be used as a basis for a model of the set of models M_k . In this way, the amount of overlap can be
20 controlled such that an even distribution of overlap is achieved.

The initial model may be processed based upon the mean model of the initial model. That is, points from the mean model of the initial model may be used as a basis for generation of further models.

25 A still further alternative is to select a set of points in an initial model associated with features of interest associated with particular image features. For example, points in regions having a high degree of curvature may be selected and for each selected point a model m may be determined as above, by selecting an area surrounding the selected point and including each point in the area in a model m .

In each of the embodiments described above, in general the size of an area modelled by a model m_{k+1} is selected such that each of the models in a set of models M_{k+1} models a

smaller region of interest than a model m_k in the corresponding set of models M_k . The size of an area modelled by a model may be chosen based upon, for example, the size of features that it is desired to be accurately modelled at stage k . As such, in some embodiments at least some models in a set of models M_{k+1} may not model a smaller region than models in a corresponding set of models M_k . For example where certain features have been shown not to benefit from further decomposition, for example during a training stage, a region selected to model the features may remain relatively large throughout the various sets of models while in other areas relatively small regions may be selected to form the basis of a model.

10

It will be appreciated that the determination of a set of models M_k is carried out in an initial stage, and as such a relatively large amount of time can be spent in selecting an effective set of models. For example, the above selection methods may each be carried out repeatedly and the efficacy of each set of models M_k that is generated may be tested against images that have been manually marked to determine how well a particular set of models M_k is able to represent variation in the test image data.

15

Whilst it has been described above that a set of models M_{k+1} generally contains a greater number of models than a set of models M_k and that each model in the set of models M_{k+1} generally model a smaller region of an anatomical region than models in the set of models M_k , it will be appreciated, that in general each set of models M_k is selected so as to effectively represent variation in the test image data. In some cases, therefore, a set of models M_{k+1} may contain the same or a smaller number of models than a set of models M_k and/or models in the set of models M_{k+1} may model regions of a similar size or a greater size than patches in a set of models M_k where such choices are able to more effectively represent variation in the test image data.

25

Referring now to Figure 5, processing to combine the fits of the models M_k of step S6 of Figure 3 is shown in further detail. At step S15 a model m in the set of models M_k is selected and at step S16 a check is performed to determine whether the fit of the selected model m is close enough to the fit of corresponding points of the combined estimate generated at step S6 for the set of models M_{k-1} . Corresponding points between models and the combined estimate are determined by the point correspondence.

Determination of whether the fit of a selected model m is close enough to the fit of corresponding points of the combined estimate generated from the set of models M_{k-1} may be determined in any convenient way. For example a geometric rule may be determined for each region of the object to be modelled. The geometric rule may 5 indicate a maximum distance a model representing the region can move relative to the starting position of the model (as indicated by the combined estimate generated from the set of models M_{k-1}) or a maximum number of degrees the model can rotate relative to the starting rotation of the model, or a combination of both. If the model does not exceed the maximum allowed, then the fit of the model is determined to be close 10 enough.

In an alternative embodiment a classification based rejection may be used in which the system is presented with a (preferably large) number of examples of a good fit of a model and a (preferably large) number of examples of a bad fit of a model in a training 15 stage. In this way, the system is able to "learn" characteristics of a good fit and a bad fit based upon residual error between fit achieved by a model and ideal fit.

If it is determined that the fit of the selected model m is close enough, then at step S17 the model m is added to a set of acceptable models A_k . If it is determined that the fit of 20 the selected model m is not close enough at step S16 or if the model m has been added to A_k at step S17, processing passes to step S18. At step S18 it is determined if there are previously unselected models in M_k . If there are previously unselected models in M_k then processing returns to step S15 where a previously unselected model is selected. Otherwise processing passes to step S19 where the models in the set A_k are combined 25 and used to update the combined estimate. Combination of models is described below with reference to Figures 6 and 7.

The processing of Figure 5 is based upon an assumption that, at each stage, the combined estimate is a reasonably accurate estimate of features in the image I. If after 30 the adjustment of the set of models M_k of step S5 points of a model m in the set of models M_k have been adjusted by greater than a predetermined amount relative to corresponding points of the combined estimate then the model m is rejected as it is assumed that the model has moved too far, and is no longer modelling the correct part

of the image. As described above, the set of models M_k overlap such that each point of the region of interest indicated in the combined estimate is modelled by a plurality of the models in M_k . Each point in the combined estimate can therefore still, in general, be adjusted based upon the set of models M_k when a model m is rejected.

5

Referring to Figure 6 the combination of models and updating of the combined estimate based upon the combination of models of step S19 of Figure 5 is shown in further detail. At step S25 a point p in the combined estimate is selected and at step S26 a set of models M_p is determined. The set of models M_p is a subset of the set of acceptable models A_k which each model the point p , determined based upon the point correspondence. At step S27 it is determined if the number of models in M_p is greater than or equal to 1. If the number of models in M_p is less than 1 then it is determined that no update can be carried out for the point p in the combined estimate and processing returns to step S25 where a point is selected that has not previously been selected.

10

If it is determined at step S27 that the number of models in M_p is greater than or equal to 1 then at step S28 a weight w_m is determined for the point p for each model m in M_p . Each weight w_m provides a weight for the contribution of a model m to the updating of point p in the combined estimate carried out at step S6 of Figure 3. Each of the models 20 in the set of models M_p may indicate a different location for point p after the model has been adjusted and a weight w_m determines how much influence a particular model m has upon the total adjustment of a point p in the combined estimate. Determination of a weight w_m for a model m is described below with reference to Figure 7. At step S29 the point p in the combined estimate is updated based upon the weighted contributions of 25 the point p in each of the models in the set M_p . That is, for a particular model point p the fit of p is updated based upon data relating to that point in each model in the set of models M_p .

25

Referring now to Figure 7 variation of a weight w_m with respect to a distance from the 30 centre of a region modelled by a model m is shown. It can be seen that at the centre of a region, indicated by circle 20 the value weight, indicated by line 21, is relatively high compared to the value weight at the edge of a region, indicated by line 22. In this way, where a point p lies relatively close to the centre of a region modelled by a model m_c ,

the contribution to the adjustment of the point p in the combined estimate based upon the model m_c will be relatively high compared to the contribution to the adjustment of the point p in the combined estimate based upon a model m_e in which the point p lies relatively close to the edge of the region modelled by the model m_e . The weights w therefore reflect the amount of information used to update a point p in any particular model m since where a point p lies relatively close to the centre of a region modelled by m_c , a relatively large amount of information surrounding the point p has been modelled by m_c and therefore used in the determination of the fit of the point p in the model m_c . It can be seen that where a point p lies at the edge of a region modelled by a model m_e , the weight w is small, but is greater than zero.

It has been described above that a model M_1 is first fitted to an image I and a set of models M_2 , is subsequently fitted to the image I based upon the fit of the model M_1 . In an alternative embodiment the fitting of the model M_1 is avoided, and the combined estimate is first established by a set of models M_i . That is, in some embodiments a plurality of models are first fitted to the image I and combined to generate a combined estimate, and subsequent sets of models M_{i+1} are fitted to the image I in the manner described above.

It has been described above that a point correspondence indicates a relationship, for each point of interest p , between points in models m which model the point of interest p and a point in a combined estimate which corresponds to the point of interest p . In an alternative embodiment the combined estimate may be a label image in which each point in the combined estimate has a label indicating a structure or feature associated with the point. For example, where it is desirable to identify an anatomical feature such as the femur in an image, some points in the combined estimate may have an associated label indicating that corresponding points in the image I are currently estimated to represent the femur. In this way, fitting a plurality of models to the image I and creation of the combined estimate can be carried out based upon labels. For example, after the set of models M_k has been fitted to the image, each model that models a particular feature contributes to updating the combined estimate such that, after combination of the models, for each point indicated as representing a particular feature in image I ,

corresponding points in the combined estimate are labelled as representing the particular feature.

Although specific embodiments of the invention have been described above, it will be
5 appreciated that various modifications can be made to the described embodiments without departing from the spirit and scope of the present invention. That is, the described embodiments are to be considered in all respects exemplary and non-limiting. In particular, where a particular form has been described for particular processing, it will be appreciated that such processing may be carried out in any suitable form
10 arranged to provide suitable output data.

CLAIMS

1. A computer-implemented method for identifying features in an image, the method comprising:

5 fitting a plurality of second models to said image, said plurality of second models together modelling a region of interest, wherein each part of said region of interest is modelled by at least two of said plurality of second models; and

identifying said features in the image based upon said fit of said plurality of second models.

10

2. A method according to claim 1, further comprising:

initially fitting a first model of said region of interest to said image;

wherein fitting the plurality of second models is based upon said fitting of said first model.

15

3. A method according to claim 2, wherein fitting said plurality of second models to said image is initialised based upon said fitting of said first model.

4. A method according to claim 2 or 3 wherein fitting said plurality of second models is limited by said fitting of said first model.

20 5. A method according to any preceding claim, wherein identifying said features in the image based upon said fit of said plurality of second models comprises:

combining data based upon said fit of said plurality of second models.

25

6. A method according to claim 5, wherein combining data based upon said fit of said plurality of second models comprises, for each part of said region of interest, applying an associated weight to data derived from each of said plurality of second models in which said part is included.

30

7. A method according to claim 6, wherein, for a particular second model, said associated weight is based upon a distance of said part of said region from a point defined with respect to the particular second model.

8. A method according to claim 7, wherein said point defined with respect to the particular second model is defined with respect to a centre of the particular model.
- 5 9. A method according to claim 5, 6, 7 or 8, wherein a particular feature is modelled by two of said second models and said combining identifies said particular feature by combining data associated with said particular feature generated by fitting each of said two of said second models.
- 10 10. A method according claim 5, 6, 7, 8 or 9, wherein said combining is based upon stored data indicating a plurality of relationships, each relationship being between at least part of a second model and at least part of at least one other second model.
- 15 11. A method according to claim 10, as dependent upon claim 2, wherein said stored data indicating a plurality of relationships indicates a relationship between at least part of a second model and part of said first model.
- 20 12. A method according to any one of claims 6 to 11, wherein said combining comprises, for each part of said region of interest included in at least two of said plurality of second models, determining a weighted average fit for said part of said region of interest based upon the fit of each of said at least two of said plurality of second models to said part of said region of interest.
- 25 13. A method according to any preceding claim, wherein identifying said features in the image based upon said fit of said plurality of second models further comprises:
determining an estimate of location of said features in said image based upon said fit of the plurality of second models.
- 30 14. A method according to claim 13, wherein determining an estimate of location of features in said image comprises:
obtaining an initial estimate of location of said features in said image; and
updating said initial estimate based upon the fit of the plurality of second models.

15. A method according to claim 14, wherein said updating said initial estimate comprises:

5 updating said initial estimate based upon said fit of one of said second models if the fit of said one of said second models satisfies a predetermined criterion.

16. A method according to claim 15, wherein said predetermined criterion is a relationship between said initial estimate of location of said features and said fit of said one of said second models.

10

17. A method according to claim 15 or 16, wherein said predetermined criterion comprises a geometric criterion.

15

18. A method according to claim 15, 16 or 17, wherein said predetermined criterion is determined for said one of said second models during a training phase, based upon a set of training images.

20

19. A method according to claim 18, wherein said predetermined criterion is based upon a plurality of data items, each data item representing a fit of said one of said second models to a respective one of said training images.

20. A method according to any one of claims 15 to 19, wherein each of said plurality of second models has an associated predetermined criterion.

25

21. A method according to any one of claims 13 to 20, further comprising:

fitting a plurality of third models to said image, said plurality of third models together modelling at least part of said region of interest, said fitting the plurality of third models being based upon said estimate of location of said features; and

30

identifying said features in the image based upon said fit of said plurality of third models.

22. A method according to claim 21, wherein a given part of said region of interest is modelled by a first number of said second models and a second different number of said third models.

5 23. A method according to claim 22, wherein said second number is greater than said first number.

24. A method according to any preceding claim, wherein said plurality of second models is an arbitrary decomposition of said object.

10

25. A method according to any preceding claim, further comprising:
generating said plurality of second models.

26. A method according to claim 25, further comprising:

15 receiving a training set of images, each image in the training set of images representing said region of interest;

wherein generating said plurality of second models of said object is based upon said training set of images.

20

27. A method according to claim 26, wherein each model in said plurality of second models is generated based upon a subset of images in said training set of images, said subset being selected by randomly sampling with replacement from said training set of images.

25

28. A method according to claim 25, 26 or 27, wherein generating said plurality of second models comprises:

generating a plurality of models, each of said plurality of models modelling at least part of said region of interest;

30 processing said plurality of models with respect to a further training set of images; and

selecting a subset of said plurality of models based upon said processing such that each part of said region of interest is represented by a plurality of models in said subset of models.

29. A method according to claim 28, wherein generating said plurality of second models comprises:

5 generating a first plurality of second models such that each part of said region of interest is modelled by a plurality of said first plurality of second models;

generating a second plurality of second models such that each part of said region of interest is modelled by a plurality of said second plurality of second models;

processing said first plurality of second models and said second plurality of second models with respect to said further training set of images; and

10 selecting one of said first plurality and said second plurality of second models of said region of interest based upon said processing.

30. A method according to claim 29, wherein processing said first plurality of second models and said second plurality of second models with respect to a training set 15 of images comprises:

fitting said first plurality of second models to each image in said further training set of images; and

fitting said second plurality of second models to each image in said further training set of images;

20 wherein selecting one of said plurality of second models of said region of interest based upon said processing is based upon said fittings.

31. A method according to any one of claims 25 to 30, wherein generating the or each plurality of second models comprises:

25 receiving data indicating a plurality of points of interest within said region of interest; and

generating the or each plurality of second models based upon said received plurality of points of interest.

30 32. A method according to claim 31, wherein generating the or each plurality of second models based upon said received data comprises:

for each of said received points of interest, selecting a part of said region of interest of predetermined size, to provide a basis for one of said plurality of models.

33. A method according to claim 31 or 32, wherein each part of said region of interest represents a feature of interest.

5 34. A method according to claim 31 or 32 as dependent upon claim 26, wherein generating said plurality of second models of said object based upon said training set of images comprises:

for each of said plurality of points, processing a corresponding part of said region of interest in each image in said training set of images.

10

35. A computer program comprising computer readable instructions configured to cause a computer to carry out a method according to any one of claims 1 to 34.

15

36. A computer readable medium carrying a computer program according to claim 35.

37. A computer apparatus for identifying features in an image, the apparatus comprising:

20

a memory storing processor readable instructions; and
a processor arranged to read and execute instructions stored in said memory; wherein said processor readable instructions comprise instructions arranged to control the computer to carry out a method according to any one of claims 1 to 34.

25

38. Apparatus for identifying features in an image, the apparatus comprising:
means for fitting a plurality of second models to said image, said plurality of second models together modelling a region of interest, wherein each part of said region of interest is modelled by at least two of said plurality of second models; and
means for identifying said features in said image based upon said fit of said plurality of second models.

30

39. A computer-implemented method for identifying features in an image, the method comprising:

fitting a plurality of second models to said image, said plurality of second models together modelling a region of interest;

fitting a plurality of third models to said image, said plurality of third models together modelling at least part of said region of interest, said fitting said plurality of third models being based upon the fit of said plurality of second models; and

identifying said features in said image based upon said fit of said plurality of third models.

40. A method according to claim 39, wherein a given part of said region of interest is modelled by a first number of said second models and a second different number of said third models.

41. A method according to claim 40, wherein said second number is greater than said first number.

42. A method according to claim 39, 40 or 41, further comprising:
initially fitting a first model of said region of interest to said image;
wherein fitting the plurality of second models is based upon said fitting of said first model.

43. A method according to claim 42, wherein fitting said plurality of second models to said image is initialised based upon said fitting of said first model and wherein fitting said plurality of third models to said image is initialised based upon said fitting of said second models.

44. A method according to any one of claims 39 to 43, wherein each part of said region of interest is represented by a plurality of said plurality of second models.

45. A computer program comprising computer readable instructions configured to cause a computer to carry out a method according to any one of claims 39 to 44.

46. A computer readable medium carrying a computer program according to claim 45.

47. A computer apparatus for identifying features in an image, the apparatus comprising:

5 a memory storing processor readable instructions; and
a processor arranged to read and execute instructions stored in said memory;
wherein said processor readable instructions comprise instructions arranged to control
the computer to carry out a method according to any one of claims 39 to 44.

48. Apparatus for identifying features in an image, the apparatus comprising:

10 means for fitting a plurality of second models to an image, said plurality of
second models together modelling a region of interest,

means for fitting a plurality of third models to said image, said plurality of third
models together modelling at least part of said region of interest, and said fitting said
plurality of third models being based upon the modified fit of said plurality of second
models; and

15 means for identifying said features in said image based upon said fit of said
plurality of third models.

49. A computer-implemented method for identifying features in an image, the
20 method comprising:

determining an estimate of location of said features in said image;
fitting a second model to part of said image; and
updating said estimate based upon said fit of said second model if the fit of said
second model satisfies a predetermined criterion.

25

50. A method according to claim 49, wherein said predetermined criterion is a
relationship between said estimate of location of said features and said fit of said second
model.

30

51. A method according to claim 49 or 50, further comprising:

fitting a first model to the image;
wherein said estimate of location of said features in the image is based upon said
fit of said first model.

52. A method according to claim 49, 50 or 51, wherein fitting said second model to part of said image is based upon said estimate of location of the features.

5 53. A method according to claim 52, wherein fitting said second model to part of said image is initialised based upon said estimate of location of the features.

54. A method according to any one of claims 49 to 53, wherein said predetermined criterion comprises a geometric criterion.

10 55. A method according to any one of claims 49 to 54, wherein said predetermined criterion is determined for said second model during a training phase, based upon a set of training images.

15 56. A method according to claim 55, wherein said predetermined criterion is based upon a plurality of data items, each data item representing a fit of said second model to a respective one of said training images.

57. A method according to any one of claims 49 to 56, wherein fitting a second model to part of said image comprises fitting a plurality of second models to said image, and wherein each of said plurality of second models has an associated predetermined criterion.

20 58. A computer program comprising computer readable instructions configured to cause a computer to carry out a method according to any one of claims 49 to 57.

25 59. A computer readable medium carrying a computer program according to claim 58.

30 60. A computer apparatus for identifying features in an image, the apparatus comprising:
a memory storing processor readable instructions; and
a processor arranged to read and execute instructions stored in said memory;

wherein said processor readable instructions comprise instructions arranged to control the computer to carry out a method according to any one of claims 49 to 57.

61. Apparatus for identifying features in an image, the apparatus comprising:
5 means for determining an estimate of location of said features in said image;
means for fitting a second model to part of said image; and
means for updating said estimate based upon said fit of said second model if the fit of said second model satisfies a predetermined criterion.

10 62. A computer-implemented method of generating a plurality of second models that together model a region of interest, the method comprising:
generating said plurality of second models such that each part of said region of interest is modelled by a plurality of said plurality of second models.

15 63. A method according to claim 62, further comprising:
receiving a training set of images, each image in the training set of images representing said region of interest;
wherein generating said plurality of second models of said object is based upon said training set of images.

20 64. A method according to claim 63, wherein each model in said plurality of second models is generated based upon a subset of images in said training set of images, said subset being selected by randomly sampling with replacement from said training set of images.

25 65. A method according to claim 62, 63 or 64, wherein generating said plurality of second models comprises:
generating a plurality of models, each of said plurality of models modelling at least part of said region of interest;
30 processing said plurality of models with respect to a further training set of images; and

selecting a subset of said plurality of models based upon said processing such that each part of said region of interest is represented by a plurality of models in said subset of models.

5 66. A method according to claim 65, wherein generating said plurality of second models comprises:

generating a first plurality of second models such that each part of said region of interest is modelled by a plurality of said first plurality of second models;

10 generating a second plurality of second models such that each part of said region of interest is modelled by a plurality of said second plurality of second models;

processing said first plurality of second models and said second plurality of second models with respect to said further training set of images; and

selecting one of said first plurality and said second plurality of second models of said region of interest based upon said processing.

15

67. A method according to claim 66, wherein processing said first plurality of second models and said second plurality of second models with respect to a training set of images comprises:

20 fitting said first plurality of second models to each image in said further training set of images; and

fitting said second plurality of second models to each image in said further training set of images;

wherein selecting one of said plurality of second models of said region of interest based upon said processing is based upon said fittings.

25

68. A method according to any one of claims 62 to 67, wherein generating the or each plurality of second models comprises:

receiving data indicating a plurality of points of interest within said region of interest; and

30 generating the or each plurality of second models based upon said received plurality of points of interest.

69. A method according to claim 68, wherein generating the or each plurality of second models based upon said received data comprises:

for each of said received points of interest, selecting a part of said region of interest of predetermined size, to provide a basis for one of said plurality of models.

5

70. A method according to claim 68 or 69, wherein each part of said region of interest of a predetermined size is centred on one of said received points of interest.

71. A method according to claim 68, 69 or 70, wherein each part of said region of interest represents a feature of interest.

72. A method according to claim 68, 69, 70 or 71 as dependent upon claim 63, wherein generating said plurality of second models of said object based upon said training set of images comprises:

15 for each of said plurality of points, processing a corresponding part of said region of interest in each image in said training set of images.

73. A computer program comprising computer readable instructions configured to cause a computer to carry out a method according to any one of claims 62 to 72.

20

74. A computer readable medium carrying a computer program according to claim 73.

75. A computer apparatus for identifying features in an image, the apparatus comprising:

25 a memory storing processor readable instructions; and
a processor arranged to read and execute instructions stored in said memory; wherein said processor readable instructions comprise instructions arranged to control the computer to carry out a method according to any one of claims 62 to 72.

30

76. Apparatus for generating a plurality of second models that together model a region of interest, the apparatus comprising:

means for generating said plurality of second models of said object such that each part of said region of interest is represented by a plurality of said plurality of second models.

5 77. A method according to any one of claims 1 to 34, 39 to 44 or 49 to 57 or 63 to 72, wherein said image or each image of said plurality of images is an image selected from the group consisting of: a computed topography image, a magnetic resonance image, a radiograph image, an x-ray image and an ultrasound image.

10 78. A method according to any one of claims 1 to 34, 39 to 44, 49 to 57, 67 or claim 77 as dependent upon any one of claims 1 to 34, 39 to 44, 49 to 57 or 67, wherein said fitting comprises applying a fitting technique selected from the group consisting of: rigid registration, non-rigid registration, active shape modelling and active appearance modelling.

15 79. A method according to any one of claims 1 to 34, 39 to 44, 49 to 57, 62 to 71, 77 or 78, wherein said models are selected from the group consisting of: statistical models, deformable templates and deformable atlases.

20 80. A method according to claim 79, wherein said models are selected from the group consisting of: shape models and appearance models.

81. A method according to claim 79, wherein said models are appearance models of a volume or of a surface.

25 82. A method according to any one of claims 1 to 34, 39 to 44, 62 to 72 or 77 to 81, wherein said plurality of second models represent an arbitrary decomposition of said object.

1/5




Fig. 1




Fig. 1A




Fig. 2

3/5




Fig. 3

4/5




Fig. 4




Fig. 5

5/5




Fig. 6




Fig. 7