



(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2018/12/12
(87) Date publication PCT/PCT Publication Date: 2019/06/20
(85) Entrée phase nationale/National Entry: 2020/06/03
(86) N° demande PCT/PCT Application No.: EP 2018/084642
(87) N° publication PCT/PCT Publication No.: 2019/115652
(30) Priorité/Priority: 2017/12/13 (GB1720791.1)

(51) Cl.Int./Int.Cl. G16H 50/20(2018.01),
A61B 8/08 (2006.01), G06T 7/00 (2017.01)

(71) Demandeur/Applicant:
OXFORD UNIVERSITY INNOVATION LIMITED, GB

(72) Inventeurs/Inventors:
UPTON, ROSS, GB;
LEESON, PAUL, GB

(74) Agent: MOFFAT & CO.

(54) Titre : PROCEDE ET APPAREIL DE MODELISATION DE DIAGNOSTIC

(54) Title: DIAGNOSTIC MODELLING METHOD AND APPARATUS

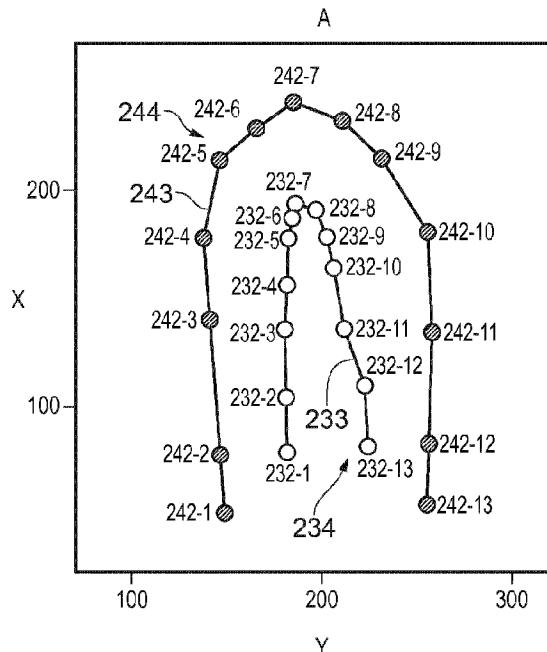


FIG. 5A

(57) Abrégé/Abstract:

The present disclosure relates to a system (100) for generating a diagnostic model. The system (100) includes a processor (108) configured to analyse a plurality of reference data sets. The reference data sets each include at least one image (230, 240). The analysis identifies at least one feature in each image (230, 240). A metric is calculated in dependence on the at least one identified feature. Outcome data associated with at least some of the reference data sets is acquired. The diagnostic model is compiled in dependence on the at least one calculated metric and the associated outcome data. The present disclosure also relates to a method of generating a diagnostic model; and a non-transitory computer-readable medium.

Abstract

The present disclosure relates to a system (100) for generating a diagnostic model. The system (100) includes a processor (108) configured to analyse a plurality of reference data sets. The reference data sets each include at least one image (230, 240). The analysis identifies at least one feature in each image (230, 240). A metric is calculated in dependence on the at least one identified feature. Outcome data associated with at least some of the reference data sets is acquired. The diagnostic model is compiled in dependence on the at least one calculated metric and the associated outcome data. The present disclosure also relates to a method of generating a diagnostic model; and a non-transitory computer-readable medium.

DIAGNOSTIC MODELLING METHOD AND APPARATUS

TECHNICAL FIELD

The present disclosure relates to a diagnostic modelling method and apparatus. More 5 particularly, but not exclusively, the present disclosure relates to a system and method for generating a diagnostic model.

BACKGROUND

Two-dimensional (2D) echocardiography is an imaging technique through which the motion of 10 the heart can be assessed under different conditions, for example resting or stress conditions. The analysis may be performed under other conditions including, for example, an intermediate stress stage and/or a recovery stage. This can highlight areas of the heart that are hypo- or dysfunctional, and can thus identify patients in which medical intervention may be necessary. A typical model of the left ventricle comprises sixteen (16) segments which are visible using 15 different 2D images of the heart. Other models of the left ventricle may, for example, comprise seventeen (17) segments. The apical inferior segment, mid inferior segment, basal inferior segment, apical anterior segment, mid anterior segment and basal anterior segment are visible in an apical two chamber image. The apical septum segment, mid septum segment, basal septum segment, apical lateral segment, mid lateral segment and basal lateral segment 20 are visible in an apical four chamber image. The anteroseptum segment, inferoseptum segment, mid inferior segment, mid anterior segment, anterolateral segment and inferolateral segment are visible in a parasternal short axis image. The apical lateral segment, the mid inferolateral segment, basal inferolateral segment, the apical septum segment, the mid septum segment, and the basal septum segment are visible in an apical three chamber image 25 (or parasternal long axis image). The behaviour of each segment can be viewed in different sections of the left ventricle. The motion of each segment of the myocardium under different conditions (such as resting and stress conditions) is currently determined by interpretation of the 2D echocardiography data by an expert cardiologist. This is performed in a categorical manner. For example, each section of the myocardial wall may be classified as having one of 30 the following reference wall motion scores: normal ("1"), hypokinetic ("2"), akinetic ("3"), dyskinetic ("4"), and unable to score ("X"). Other classifications may be used, for example defining five (5) or seven (7) discrete scores for each segment. The known techniques rely on subjective classification and may prove a time consuming exercise.

35 It has been recognised that image quantification tools need to allow for the following: (i) changing underlying disease pathophysiology over time; (ii) disease pathology variations with

geographical location and changing nature of the patient population being referred for the test; and (iii) the changing understanding of what is defined as disease or what is disease causing.

At least in certain embodiments, the present invention seeks to provide an improved method

5 and apparatus for generating a diagnostic model.

SUMMARY OF THE INVENTION

Aspects of the present invention relate to a system for generating a diagnostic model, a method of generating a diagnostic model, and a non-transitory computer-readable medium as

10 claimed in the appended claims.

According to a further aspect of the present invention there is provided a system for generating a diagnostic model, the system comprising a processor configured to:

analyse a plurality of reference data sets, each reference data set comprising at least

15 one image, the analysis comprising identifying at least one feature in each image;

calculate at least one metric in dependence on the at least one identified feature;

acquire outcome data associated with at least some of the reference data sets; and

compile the diagnostic model in dependence on the at least one calculated metric and the associated outcome data. The use of outcome data establishes a feedback loop which can be used to refine and develop the diagnostic model so as to reflect different outcome scenarios. The outcome data may be used in machine learning algorithms to adjust thresholds and/or weightings. The outcome data may potentially facilitate identification of new biomarkers for diagnostics. The outcome data may be used to modify or adapt the algorithms used to generate the diagnostic model.

25

At least in certain embodiments, the reference data sets may each comprise a plurality of images. Each reference data set may comprise a first image and a second image. The processor may be configured to analyse each reference data set to identify at least one first

30 feature in the first image, and identify at least one second feature in the second image, each at least one first feature being paired with a corresponding one of the at least one second feature.

The processor may be configured to analyse each reference data set to compare each pair of corresponding first and second features to identify one or more difference therebetween. Each pair of corresponding first and second features relate to the same feature in both the first and second images. The first and second images may relate to the same

35 region, but may be acquired at different times.

The one or more difference identified between the first and second features may, for example, comprise one or more of the following: opacity, brightness, contrast, cross-sectional area, size (in one dimension, two dimensions or three dimensions), position (in one dimension, two dimensions or three dimensions) and orientation (about one axis, two axes, or three axes).

5

The processor may be configured to calculate the at least one metric in dependence on the one or more difference identified between each pair of corresponding first and second features.

10 The system may be suitable for generating a diagnostic model for diagnosing a heart condition. The first image may comprise a first end systolic image and the second image comprises a second end diastolic image. The one or more difference identified between each pair of corresponding first and second features may represent a cardiac cyclic change. The at least one metric may be calculated in dependence on the identified cardiac cyclic change in

15 respect of each reference data set.

The processor may be configured to label the images in dependence on the outcome data. The labels may distinguish between different classifications. Each label may, for example comprise a classification indicating a presence or an absence of a condition or a disease.

20 Each label may comprise a grade indicating a severity of a condition, for example comprising an indication of the severity of stenosis.

The outcome data may comprise diagnostic information. The diagnostic information may comprise a record of a cardiac event, such as a myocardial infarction.

25

The outcome data may comprise a record of stenosis greater than a threshold value. The threshold value may be defined as a percentage of stenosis. The threshold may, for example, be defined as 60%, 70%, 80% or 90%.

30 The outcome data may be generated after acquisition of the echocardiograph images. For example, the outcome data may be generated 6 months, 12 months, 18 months, 24 months or longer after acquisition of the echocardiograph images.

The processor may be configured to update the diagnostic model when the outcome data becomes available or when the outcome data is updated.

The processor may be configured to determine a weighting for at least some of the calculated metrics. Each weighting may be determined in dependence on the outcome data associated with a given one of the reference data sets.

5 The diagnostic model may be compiled in dependence on the calculated metrics and the associated weightings.

The processor may be configured to analyse further reference data sets. The further reference data sets may each comprise at least one image. The processor may be configured to update

10 the diagnostic model in dependence on the analysis of the further reference data sets.

The at least one image may comprise an ultrasound image. The at least one image may comprise an echocardiograph image. Alternatively, or in addition, the at least one image may comprise a magnetic resonance image (MRI) or a computed axial tomography (CAT) image.

15

According to a further aspect of the present invention there is provided a method of generating a diagnostic model, the method comprising:

analysing a plurality of reference data sets, each reference data set comprising at least one image, the analysis comprising identifying at least one feature in each image;

20

calculating at least one metric in dependence on the at least one identified feature;

acquiring outcome data associated with at least some of the reference data sets; and

compiling the diagnostic model in dependence on the calculated metrics and the associated outcome data.

25

At least in certain embodiments, the reference data sets may each comprise a plurality of images. Each reference data set may comprise first and second images. The method may comprise analysing each reference data set to identify at least one first feature in the first image, and identify at least one second feature in the second image, each at least one first feature being paired with a corresponding one of the at least one second feature. The method

30

may comprise comparing each pair of corresponding first and second features to identify one or more difference therebetween. Each pair of corresponding first and second features relate to the same feature identified in both the first and second images. The first and second images may relate to the same region, but may be acquired at different times. The one or more difference identified between the first and second features may, for example, comprise one or

35

more of the following: opacity, brightness, contrast, cross-sectional area, size (in one dimension, two dimensions or three dimensions), position (in one dimension, two dimensions or three dimensions) and orientation (about one axis, two axes, or three axes).

The method may comprise calculating the at least one metric in dependence on the one or more difference identified between each pair of corresponding first and second features.

5 The system may be suitable for generating a diagnostic model for diagnosing a heart condition. The first image comprises a first end systolic image and the second image comprises a second end diastolic image. The one or more difference identified between each pair of corresponding first and second features represents a cardiac cyclic change.

10 The outcome data may be used to label the corresponding reference data sets. The labels may distinguish between different classifications. Each label may, for example comprise a classification indicating a presence or an absence of a condition or a disease. Each label may comprise a grade indicating a severity of a condition, for example comprising an indication of the severity of stenosis.

15

The outcome data may comprise diagnostic information. The diagnostic information may comprise a record of a cardiac event, such as a myocardial infarction.

The outcome data may comprise a record of stenosis greater than a threshold value.

20

The outcome data may be generated after acquisition of the echocardiograph images. For example, the outcome data may be generated 6 months, 12 months, 18 months, 24 months or longer after acquisition of the echocardiograph images.

25 The method may comprise updating the diagnostic model when the outcome data is updated and/or when new outcome data becomes available.

30 The method may comprise determining a weighting for at least some of the calculated metrics. Each weighting may be determined in dependence on the outcome data associated with a given one of the reference data sets. The diagnostic model may be compiled in dependence on the calculated metrics and the associated weightings.

35 The method may comprise adding further reference data sets of incrementally. The method may comprise analysing the further reference data sets. The method may comprise updating the diagnostic model in dependence on the analysis of the further reference data sets.

The at least one image may comprise an ultrasound image. The at least one image may comprise an echocardiograph image. Alternatively, or in addition, the at least one image may comprise a magnetic resonance image (MRI) or a computed axial tomography (CAT) image.

5 According to a further aspect of the present invention there is provided a system for generating a diagnostic model for diagnosing a heart condition, the system comprising a processor configured to:

analyse a plurality of reference data sets, each reference data set comprising at least first and second echocardiograph images, the analysis comprising identifying an end systolic

10 image and an end diastolic image within each reference data set;

compare the end systolic image and the end diastolic image in each reference data set to identify a cardiac cyclic change;

calculate at least one metric in dependence on the identified cardiac cyclic change in respect of each reference data set;

15 acquire outcome data associated with at least some of the reference data sets; and

compile the diagnostic model in dependence on the calculated metrics and the associated outcome data.

The outcome data may be acquired a period of time after the at least one image. For example,

20 the outcome data may be generated 6 months, 12 months, 18 months, 24 months or longer after acquisition of the at least one image.

According to an aspect of the present invention there is provided a method of generating a diagnostic model for diagnosing a heart condition, the method comprising:

25 analysing a plurality of reference data sets, each reference data set comprising at least first and second echocardiograph images, the analysis comprising identifying an end systolic image and an end diastolic image within each reference data set;

comparing the end systolic image and the end diastolic image in each reference data set to identify a cardiac cyclic change;

30 in respect of each reference data set, calculate at least one metric in dependence on the identified cardiac cyclic change;

acquiring outcome data associated with at least some of the reference data sets; and

compiling the diagnostic model in dependence on the calculated metrics and the associated outcome data.

According to a further aspect of the present invention there is provided a non-transitory computer-readable medium having a set of instructions stored therein which, when executed, cause a processor to perform the method described herein.

5 Any control unit or controller described herein may suitably comprise a computational device having one or more electronic processors. The system may comprise a single control unit or electronic controller or alternatively different functions of the controller may be embodied in, or hosted in, different control units or controllers. As used herein the term "controller" or "control unit" will be understood to include both a single control unit or controller and a plurality
10 of control units or controllers collectively operating to provide any stated control functionality. To configure a controller or control unit, a suitable set of instructions may be provided which, when executed, cause said control unit or computational device to implement the control techniques specified herein. The set of instructions may suitably be embedded in said one or more electronic processors. Alternatively, the set of instructions may be provided as software
15 saved on one or more memory associated with said controller to be executed on said computational device. The control unit or controller may be implemented in software run on one or more processors. One or more other control unit or controller may be implemented in software run on one or more processors, optionally the same one or more processors as the first controller. Other suitable arrangements may also be used.

20

Within the scope of this application it is expressly intended that the various aspects, embodiments, examples and alternatives set out in the preceding paragraphs, in the claims and/or in the following description and drawings, and in particular the individual features thereof, may be taken independently or in any combination. That is, all embodiments and/or
25 features of any embodiment can be combined in any way and/or combination, unless such features are incompatible. The applicant reserves the right to change any originally filed claim or file any new claim accordingly, including the right to amend any originally filed claim to depend from and/or incorporate any feature of any other claim although not originally claimed in that manner.

30

BRIEF DESCRIPTION OF THE DRAWINGS

One or more embodiments of the present invention will now be described, by way of example only, with reference to the accompanying figures, in which:

35 Figure 1 is a schematic view of an echocardiography system according to an embodiment of the invention;

Figure 2 shows schematically a four-chamber view of a heart;

Figure 3 is a flow diagram showing the main steps of a diagnostic method performed by the system of Figure 1;

Figure 4A shows an end systole image captured by the echocardiography system shown in Figure 1;

5 Figure 4B shows an end diastole image captured by the echocardiography system shown in Figure 1;

Figure 5A shows contour data sets composed of end systole contour points and end diastole contour points;

10 Figure 5B shows a plurality of elements generated from the contour data sets shown in Figure 5A;

Figure 6A illustrates generation of an area metric for each element illustrated in Figure 5B;

Figure 6B illustrates generation of a mean distance metric for each element illustrated in Figure 5B;

15 Figure 6C illustrates generation of a rectangularity metric for each element illustrated in Figure 5B;

Figure 6D illustrates generation of a solidity metric for each element illustrated in Figure 5B;

Figure 7A illustrates the analysis of the area metric for a rest condition;

20 Figure 7B illustrates the analysis of the mean distance metric for a rest condition;

Figure 7C illustrates the analysis of the rectangularity metric for a rest condition;

Figure 7D illustrates the analysis of the solidity metric for a rest condition;

Figure 8A illustrates the analysis of the area metric for a stress condition;

Figure 8B illustrates the analysis of the distance metric for a stress condition;

25 Figure 8C illustrates the analysis of the rectangularity metric for a stress condition;

Figure 8D illustrates the analysis of the solidity metric for a stress condition;

Figure 9A illustrates a reference data model based on bivariate analysis of the normalised area and mean distance metrics;

30 Figure 9B shows a normally distributed data set applied to the reference data model illustrated in Figure 9A;

Figure 10A illustrates the multivariate analysis of each of the metrics for a first rest condition in a two-chamber apical image;

Figure 10B illustrates the multivariate analysis of each of the metrics for a second rest condition in a four-chamber apical image;

35 Figure 11A illustrates the multivariate analysis of each of the metrics for a first stress condition in a two-chamber apical image;

Figure 11B illustrates the multivariate analysis of each of the metrics for a second stress condition in a four-chamber apical image;

Figure 12A shows scoring applied to elements in a rest condition;

Figure 12B shows scoring applied to elements in a stress condition; and

5 Figure 13 shows a decision tree composed of a series of decision points defining threshold values.

DETAILED DESCRIPTION

An echocardiography system 100 in accordance with an embodiment of the present invention 10 will now be described with reference to the accompanying figures. The echocardiography system 100 is operable to analyse images of a heart 200 and to score the cardiac cyclic motion.

As shown in Figure 1, the echocardiography system 100 comprises a transducer array 102 15 arranged to be located close to the body of a patient 104, typically as close to the heart as possible, a processing unit 106 which includes a processor 108 which may be a digital electronic processor, a memory 110 such as a hard disk, and a display 112, such as a flat screen monitor or LED display. The system may further include a user input device, for example a touchscreen 114 integrated into the display 112, which provides a user input 20 allowing a user to provide inputs to the echocardiography system 100. Other user inputs such as a mouse, touchpad or keyboard may of course be used. The processing unit 106 is connected to the transducer array 102 and is arranged to control the transducer array as a phased array so as to emit an ultrasound beam which scans across the patient in a series of pulses, and detect reflected ultrasound from the heart from each pulse. One scan of the heart 25 builds up a single image, and the scan is repeated at typically 25 to 50 images per second to build up a real time video image of the heart showing its movement during the cardiac cycle. Each image may be stored in the memory 110 as an image data set which may comprise, for example, intensity values for each of the pixels of which the image is made up. While the system is described herein in general terms, suitable echocardiography systems include, for 30 example the Philips Epic iE33, GE vivid e9, or portable systems such as the Philips CX50, or hand-held systems.

The process of echocardiography is well known and is not described herein in detail. There are several different imaging methods, but the echocardiography system 100 in accordance 35 with the present embodiment uses two-dimensional imaging. It is known to provide images on several different planes through the heart, which show different aspects of the four main chambers of the heart, the left ventricle (LV), right ventricle (RV), left atrium (LA) and right

atrium (RA). Such views include, for example, an apical four chamber view, an apical two chamber view, an apical three chamber view and parasternal long and short axis views. In each case, while a single still image can be obtained, typically a series of views is acquired over the cycle of the heart so that its movement can be recorded and analysed. The 5 echocardiography system 100 may utilise one or more of the aforementioned views to score the cardiac cyclic motion of the heart 200.

A four-chamber apical image of a heart 200 is shown in Figure 2 by way of example. The image comprises a 2D plane of the heart 200. The image shows a left ventricle (LV) 202, a 10 right ventricle (RV) 204, a left atrium 206, a right atrium 208 and a septum 210. An apex 212, a lateral wall 214, a base 216 and an inner wall 218 of the left ventricle 202 are also visible in the four-chamber apical view. A longitudinal axis 220 of the left ventricle 202 extends through the apex 212. The left ventricle 202 has first and second sides 222, 224 disposed on opposing sides of the longitudinal axis 220.

15

The processing unit 106 analyses the four-chamber apical image to implement the scoring techniques described herein. Alternatively, or in addition, the processing unit 106 may utilise one or more of the following: a two-chamber apical image, a parasternal short axis image and a three-chamber apical view. Other echocardiograph images could be used by the processing 20 unit 106 to implement the scoring techniques described herein. The processing unit 106 may use various combinations of the echocardiograph images provide scoring for the sixteen (16) segments of the left ventricle. The processing unit 106 may analyse a plurality of images and score the cardiac cyclic motion in dependence on the metrics for multiple images. The processing unit 106 may qualitatively assess the available images and prioritise an image 25 determined as providing a clearer representation of the cardiac cyclic motion of a particular section of the heart wall. A Cartesian coordinate system is defined comprising a vertical axis (referred to as the y axis herein) extending through the apex 212 of the left ventricle 202 and extending along its longitudinal axis, and a horizontal axis (referred to as the x axis herein) through the mid-point of the left ventricle 202 half way between the apex 212 and the base 30 216.

A block diagram representing operation of the echocardiography system 100 is shown in Figure 3. The echocardiography system 100 is arranged to acquire a sequence of 2D images and store them in memory 110 (BLOCK 300). The images may be acquired over a single 35 cardiac cycle, and may include for example between ten (10) and fifty (50) images covering one cycle. The echocardiography system 100 may perform a single scan or more than one scan. For example, the echocardiography system 100 may perform first and second scans.

The first scan may be performed when the patient is under rest conditions and the second scan may be performed when the patient is under stress condition. The echocardiography system 100 may optionally perform one or more intermediate scan between the rest condition and the stress condition, for example during a recovery phase as heart rate returns to normal

5 after being stressed. The acquisition of the images can be carried out on a conventional echocardiography system. The subsequent analysis of the images can be carried out using the same processing unit 106 that forms part of the echocardiography system as shown in Figure 1. However, the images may be downloaded onto a computer, such as a laptop or PC, which has a processor, memory, user input and display, which operate for this purpose in the

10 same way as those of the processing unit 106, and the further analysis of the images may be carried out on that computer under the control of dedicated software. It will be understood that the images may be retrieved from a PACS (picture archiving and communication system). Alternatively, or in addition, images may be transmitted to an external server for processing. The images may be anonymised prior to transmission.

15

The image closest to end systole, i.e. maximum contraction during the cardiac cycle, and the image closest to end diastole, i.e. maximum volume during the cardiac cycle, are identified for the left ventricle 202 (BLOCK 302). This can be done by a user viewing the images on the display 112 and selecting a first image 230 as closest to end systole (referred to herein as the end systole image 230), and a second image 240 as closest to end diastole (referred to herein as the end diastole image 240). The end systole image 230 and the end diastole image 240 are acquired at first and second times respectively in the cardiac cycle. An exemplary end systole image 230 is shown in Figure 4A, and an exemplary end diastole image 240 is shown in Figure 4B. The selection of the end systole image 230 and the end diastole image 240 may be made by the user on the basis of an assessment and comparison of the volume of the left ventricle 202 in each of the images as judged by eye, or by noting the points of opening and closing of the mitral valve, or using the QRS complex on an ECG plot, or by any combination of these. Alternatively, the processor 108 may be arranged to use image processing techniques to identify the end systole image 230 and the end diastole image 240. The image processing techniques may, for example, determine the volume of the left ventricle 202 in each of the images. The processor may identify the image with the smallest left ventricle volume as the end systole image 230; and the image with the largest left ventricle volume as the end diastole image 240. Alternatively, the image processing techniques may identify and track movements of image elements which are persistent across multiple images to identify the end systole image 230 and the end diastole image 240. The reversal in the direction of movement of the persistent image elements may be used to identify end systole and end diastole, for example. The end systole image 230 and the end diastole image 240 are identified

in the memory 110, for example being marked with an appropriate flag, so that they can be selected and viewed by a user.

The inner wall 218 of the left ventricle 202 is identified at end systole in the end systole image 230, and at end diastole in the end diastole image 240 (BLOCK 304). The left ventricle 202 is contoured (or mapped) at end diastole in the end systole image 230 and at end systole in the end diastole image 240 (BLOCK 306). The contouring of the left ventricle 202 comprises identifying a plurality of end systole contour points 232-n around the inner wall 218 in the end systole image 230; and a plurality of end diastole contour points 242-n around the inner wall 218 in the end diastole image 240. A first continuous curve is plotted between the end systole contour points 232-n to form an end systole contour line 233; and a second continuous curve is plotted between the end systole contour points 242-n to form an end diastole contour line 243. The end systole contour line 233 and the end diastole contour line 243 may comprise straight lines and/or curved lines. The end systole contour line 233 and the end diastole contour line 243 may, for example, be profiled to match a boundary identified in the end systole image 230 and the end diastole image 240 respectively.

The end systole contour points 232-n and the end systole contour line 233 form an end systole contour data set 234; and the end diastole contour points 242-n and the end diastole contour line 243 form an end diastole contour data set 244. Each end systole contour point 232-n in the end systole contour data set 234 is paired with a corresponding one of the end diastole contour points 242-n in the end diastole contour data set 244. The resulting pairs of end systole and end diastole contour points 232-n, 242-n represent changes in the motion of the wall of the heart 200 during a cardiac cycle. The pairs of end systole and end diastole contour points 232-n, 242-n may correspond to the same feature of the left ventricle 202, albeit in different locations in the end systole image 230 and the end diastole image 240 due to the wall motion during the cardiac cycle. In the present embodiment, thirteen (13) end systole and end diastole contour points are identified in the end systole image 230 and the end diastole image 240. The end diastolic contour points and the end systolic contour points are labelled 1 to 13 according to their position along the endocardium (i.e. n=1, 2, 3, ...13). The end systole and end diastole contour data sets 234, 244 are combined, as shown in Figure 5A.

As shown in Figure 5B, the end systole contour points 232-n and the end diastole contour points 242-n form a plurality of elements E_i (where i is a whole number) representing the cardiac cyclic motion of the internal wall (BLOCK 308). As described herein, the elements E_i are analysed and scored to grade the cardiac cyclic motion of the corresponding section of the inner wall 218 of the left ventricle 202. In the illustrated arrangement, three (3) pairs of end

systole contour points 232-n and end diastole contour points 242-n (i.e. three (3) end systole contour points 232-n and three (3) end diastole contour points 242-n from the respective end systole and end diastole contour data sets 234, 244) define each element E_i . In the illustrated example, the end systole image 230 and the end diastole image 240 are two-chamber apical

5 images. The changes in the cardiac cyclic motion of the internal wall are represented by six (6) elements E_{1-6} . Each of the elements E_{1-6} is in the form of a planar (two-dimensional) polygon. The elements E_{1-6} correspond to a respective segment of the model of the left ventricle. In particular, a first element E_1 corresponds to the basal inferior segment; a second element E_2 corresponds to the mid inferior segment; a third element E_3 corresponds to the 10 apical inferior segment; a fourth element E_4 corresponds to the apical anterior segment; a fifth element E_5 corresponds to the mid anterior segment; and a sixth element E_6 corresponds to the basal anterior segment. It will be understood that the cardiac cyclic changes may be represented by a different number of elements E_i , for example less than six (6) elements or more than six (6) elements.

15

The elements E_i are analysed to generate at least one wall motion metric for scoring (i.e. classifying or grading) the cardiac cyclic motion of the corresponding sections of the heart 200 (BLOCK 310). The analysis of the elements E_i is described in more detail herein. The generated metric is compared to a predefined reference data model to score the wall motion 20 (BLOCK 312). The results of the scoring are then output, for example to a screen or display (BLOCK 314). The scoring may be reviewed by a clinician.

The contouring of the left ventricle 202 will now be described in more detail. The contouring may be performed by an echocardiographer; or using suitable image processing techniques.

25 Echo images of a left ventricle 202 acquired with a contrast agent are shown in Figures 4A and 4B. The end systole image 230 is shown in Figure 4A; and the end diastole image 240 is shown in Figure 4B. The apex 212 of the left ventricle 202 can be located as the top of the left ventricle 202, and the base 216 of each side 222, 224 can be located from the shape of the inner wall 218. The longitudinal (Y) axis is defined as the reference line passing through the 30 apex 212 and the midpoint between the base of the two sides 222, 224. The x axis can then be defined as the line perpendicular to the y axis half way between the apex and the midpoint between the two sides of the base 216. The mid-point on each side 222, 224 can be identified as the point where the x axis intersects the side wall on that side 222, 224. The intermediate end systole contour points 232-n and the end diastole contour points 242-n may be identified 35 by subdividing the regions between the apex 212 and the mid-point on each side 222, 224; and by subdividing the region between the mid-point and the base on each side 222, 224.

As mentioned above, each of these end systole contour points 232-n and the end diastole contour points 242-n may be identified by a user. Alternatively, image processing may be used to identify the end systole contour points 232-n and the end diastole contour points 242-n. If image processing is used, the outline of the left ventricle 202 is first identified as the boundary

5 between the lighter area within the left ventricle 202 and the darker area of the myocardium forming the walls around it (or vice versa for images acquired without use of a contrast agent). Suitable algorithms for identifying such boundaries are well known. Once the boundary has been identified, the algorithm may then be arranged to identify the highest point (maximum y value) of the boundary as being the apex 212, and the points where the boundary changes 10 direction at the lower end as the base 216. Again, algorithms for analysing the radius and direction of curvature, and how that changes around the boundary, can be used to identify these points, and the points at the lower end of the apex 212. The coordinates of each of the end systole contour points 232-n and the end diastole contour points 242-n are determined 15 with reference to the coordinate system. The scale of the images acquired by the echocardiography system 100 is known. Thus, the coordinates of each of the end systole contour points 232-n and the end diastole contour points 242-n define the position of the point in the plane of the corresponding image. The distance between the contour points in each pair indicates the distance moved by the corresponding section of the heart 200 between end systole and end diastole.

20

The analysis of the elements E_i to generate wall motion metrics will now be described with referenced to Figures 6A to 6D. As shown in Figures 6A, each element E_i is in the form of a polygon having n sides. As shown in Figure 6A, an element area A of each element E_i is calculated from the first and second sets of contour data 234, 244 by means of a shoelace 25 formula:

$$A = \frac{1}{2} \left| \sum_{i=1}^{n-1} x_i y_{i+1} + x_n y_1 - \sum_{i=1}^{n-1} x_{i+1} y_i - x_1 y_n \right| \quad (\text{Equation 1})$$

30 where n = the number of sides of the polygon; and

(x_i, y_i) = the vertices of the polygon ($i = 1, 2, \dots, 6$).

The calculated area A of each element E_i is then normalised as a fraction of the total area represented by the total area of the end-diastolic contour points.

35

As shown in Figure 6B, the Euclidean distance (d) between each pair of end-diastolic and end-systolic end systole contour points 232-n and the end diastole contour points 242-n is computed using the equation:

$$5 \quad d = \sqrt{\sum_{i=1}^n (q_i - p_i)^2} \quad (\text{Equation 2})$$

where n = the number of dimensions;
 p = the co-ordinates of the end diastolic contour point; and
10 q = the co-ordinates of the end systolic contour point.

The mean distance (d) for each element E_i is then calculated (i.e., $\bar{x}(d_1, d_2, d_3)$ for the first element E_1 , $\bar{x}(d_3, d_4, d_5)$ for the second element E_2 , and so on). The mean distance is subsequently normalised as a fraction of the total perimeter distance of the end diastolic
15 contour points.

As shown in Figure 6C, a rectangularity of each element E_i was calculated as the ratio between the area of each element (A_i) and the area of its minimum bounding rectangle (R_i):

$$20 \quad \text{Rectangularity}_i = \frac{A_i}{R_i} \quad (\text{Equation 3})$$

where A_i = area of each element E_i ; and
 R_i = area of the minimum bounding rectangle.

25 As shown in Figure 6D, a solidity S_i of each element E_i was calculated as the ratio between the element's area (A_i) and the area of its convex hull (H_i):

$$30 \quad S_i = \frac{A_i}{H_i} \quad (\text{Equation 4})$$

where A_i = area of each element E_i ; and
 H_i = area of the corresponding convex hull.

In order to assess the correlation between the calculated metrics and the reference wall motion scores, a reference data set comprising raw (i.e. unprocessed) two-dimensional echocardiography data was analysed. The reference data set was composed of historic data

5 comprising end diastolic images and end systolic images for a group of patients. The end diastolic images and the end systolic images were analysed in accordance with the techniques described herein to identify the end systole contour points 232-n and the end diastole contour points 242-n. Elements E_i corresponding to respective segments of a standard model of the left ventricle 202 were thereby identified. The elements E_i were analysed using the techniques

10 described herein to calculate the following metrics: normalised area A , normalised mean distance d , rectangularity, and solidity S_i . The metrics were generated for rest and stress conditions for each element E_i . The elements E_i were also independently scored by two cardiologists using a standard scoring system consisting of the reference wall motion scores: normal ("1"), hypokinetic ("2"), akinetic ("3"), dyskinetic ("4"), and unable to score ("X"). Any

15 scores that were discrepant between the two reference data sets were reviewed and a consensus reached. Elements with a wall motion score of "X" were removed from the reference data set ($n = 2$). Due to the low number of elements in the available reference data set having a wall motion score of "4" ($n = 2$), these were also removed from the analysis. Thus, in the present embodiment, each element E_i from the reference data set was scored as normal

20 ("1"), hypokinetic ("2"), akinetic ("3"). The analysis was repeated for a rest condition and a stress condition for each patient. As described herein, the metrics calculated through analysis of the raw reference data set are used to generate a reference data model against which the calculated metrics may be compared. The reference data model is generated for each element E_i . The reference data model may be a univariate model or a multivariate model. The reference

25 data model may be stored in the memory 110 of the echocardiography system 100. Alternatively, the reference data set may be stored in the memory 110 and the reference data model generated by the processing unit 106. This approach may enable dynamic comparisons, for example in respect of particular metrics or combinations of metrics.

30 A plot of the calculated metrics for each element E_i and the wall motion score allocated by the cardiologists (i.e. normal "1", hypokinetic "2", akinetic "3") for the rest condition are shown in Figures 7A-7D. The second quartile of the metrics is represented for each reference wall motion score by a box plot comprising a median line for that set of metrics. The normalised element area A for each reference wall motion score is shown in Figure 7A. The normalised

35 mean distances d for each reference wall motion score is shown in Figure 7B. The calculated solidity for each reference wall motion score is shown in Figure 7C. The calculated rectangularity for each reference wall motion score is shown in Figure 7D. A Wilks-Lambda

non-parametric, multivariate test statistic of $P < 0.05$ is determined in each wall motion score group. A statistically significant correlation is identified between the calculated metrics and the allocated wall motion score.

5 A plot of the calculated metrics for each element E_i and the wall motion score allocated by the cardiologists (i.e. normal “1”, hypokinetic “2”, akinetic “3”) for the stress condition are shown in Figures 8A-8D. The second quartile of the metrics is represented for each reference wall motion score by a box plot comprising a median line for that set of metrics. The normalised element area A for each reference wall motion score is shown in Figure 8A. The normalised
10 mean distances d for each reference wall motion score is shown in Figure 8B. The calculated solidity for each reference wall motion score is shown in Figure 8C. The calculated rectangularity for each reference wall motion score is shown in Figure 8D. A Wilks-Lambda non-parametric, multivariate test statistic of $P < 0.05$ is determined in each wall motion score group. A statistically significant correlation is identified between the calculated metrics and the
15 allocated wall motion score.

The scoring for each element E_i can be calculated in dependence on one of the calculated metrics. The processing unit 106 may be configured to define a univariate distribution, for example a univariate normal distribution. By way of example, the scoring can correspond to a
20 z-score (standard score) for one of the calculated metrics. The z-score indicates how many standard deviations a calculated metric is from the population mean in units of standard deviation. The processing unit 106 may be configured to allocate a score to each element E_i corresponding to the determined z-score. However, the accuracy of the score calculated for each element E_i may be improved referencing two or more of the calculated metrics. The
25 processing unit 106 may be configured to define a multivariate distribution, for example a multivariate normal distribution. The processing unit 106 may be configured to define a bivariate distribution or a higher dimensional distribution. The processing unit 106 may be configured to calculate a ‘distance’ of the calculated metric from a reference population. This technique enables analysis to be performed in higher dimensions. One approach is to use the
30 distance from the first principal component of the data PC1. Alternatively, or in addition, the Mahalanobis distance may be calculated by the processing unit 106. Other statistical analysis techniques are also appropriate.

By determining the correlation between the mean distance and the area of each element E_i , a
35 score can be determined for the wall section corresponding to each element E_i . In the present embodiment, the scoring comprises a continuous scale, rather than the traditional scoring system which relies on discrete values. A mock representation of the correlation between z-

scaled element areas and mean distances is illustrated in Figure 9A. A lower left quadrant Q3 represents those elements E_i identified as having potentially abnormal wall motion. A set of three (3) ellipses E1 to E3 represent the confidence intervals (CI) for the distribution of the data points: the inner ellipse E1 represents CI=68%, the middle ellipse E2 represents CI=95% and the outer ellipse E3 represents CI=99%. A centroid O of the data is shown; and a line PC1 represents an orthogonal regression line through the data (i.e., the first principal component of the data PC1). A set of markers M1-M3 are representative of data points which are being scored. A set of randomly generated, normally distributed data with a covariance of 0.56 ($n = 1,000$) is illustrated in Figure 9B. Each data point is coloured according to the continuous scoring determined in accordance with the analysis techniques described herein.

The processing unit 106 in accordance with the present embodiment implements a continuum approach for scoring each element E_i . The principal component models are constructed in dependence on the z-scaled metrics of each element E_i , as described herein. This is performed for each elements E_i derived from the end systole image 230 and the end diastole image 240. The description herein focuses on the six (6) elements E_i corresponding to the segments visible in the standard model of the two-chamber apical images. It will be understood that the same techniques may be implemented in respect of additional elements E_i corresponding to other segments of the left ventricle 202, for example by analysis of three-chamber apical images and/or four-chamber apical images. The analysis is performed independently in respect of end systole and end diastole images 230, 240 acquired for rest and stress conditions. The processing unit 106 may compare the results of the analysis in respect of the rest and stress conditions.

The scoring of the elements E_i in dependence on a bivariate analysis based on two calculated metrics is visualised in Figures 9A and 9B. The metrics in the present case are the normalised area and the mean distance of each element E_i . A reference data model is generated in dependence on the normalised area and the mean distance of the elements E_i identified through analysis of the reference data set. In the present case, only those elements with an allocated wall motion score of "1" were included in the generation of the reference data model. In order to generate a score for a given element E_i , the processing unit 106 calculates the corresponding metrics for that element E_i . The processing unit 106 calculates the normalised area and the mean distance of elements E_i identified through analysis of the end systole image 230 and the end diastole image 240 for a patient. The implementation described herein with reference to Figures 9A and 9B utilises bivariate analysis based on the normalised area and the mean distance of each element E_i . It will be understood that other combinations of the metrics may be used for scoring each element E_i . For example, the bivariate analysis may

combine the mean distance and solidity metrics; or the normalised area and rectangularity metrics.

The processing unit 106 may be configured to perform multivariate analysis. The processing

5 unit 106 may be configured to combine each of the metrics described herein, namely: the normalised area A , the normalised mean distance d , the rectangularity, and the solidity S_i . The score for each element E_i may be calculated in dependence on the multivariate analysis of the four (4) calculated metrics. Plots of the score calculated in dependence on a first principal component PC1 and the allocated wall motion score (i.e. normal “1”, hypokinetic “2”, akinetic 10 “3”) are shown in Figures 10, 10B, 11A and 11B. A plot of the score calculated in dependence on a first principal component PC1 of the multivariate analysis of a rest condition in a two-chamber apical image is illustrated in Figure 10A. A plot of the score calculated in dependence on a first principal component PC1 of the multivariate analysis of a rest condition in a four-chamber apical image is illustrated in Figure 10B. A plot of the score calculated in dependence 15 on a first principal component PC1 of the multivariate analysis of a stress condition in a two-chamber apical image is illustrated in Figure 11A. A plot of the score calculated in dependence on a first principal component PC1 of the multivariate analysis of a stress condition in a four-chamber apical image is illustrated in Figure 11B.

20 The calculated metrics for each element E_i are compared to the reference data model for a corresponding element E_i . The score for each element E_i is calculated in dependence on this comparison. The score represents a value of the first principal component of the new data; i.e. how far the calculated metrics are from the centroid of the data and thus how different they are from the reference data. The score can be calculated on a continuous scale. An example 25 of this can be seen in Figures 12A and 12B where each element has been shaded according to the continuous scoring scale described herein. The scored images shown in Figure 12A represent two chamber data for a rest condition; and the scored images shown in Figure 12B represent two chamber data for a stress condition.

30 The processing unit 106 may be configured to calculate different metrics for scoring each element E_i . These metrics may be used in addition to, or instead of the metrics described herein for the univariate and multivariate analysis. The processing unit 106 may, for example, calculate one or more of the following metrics: shear; strain; coefficient of variation of the distances in an element; and aspect ratio. The processing unit 106 may also calculate a 35 distance metric other than the distance between the pairs of points described herein. For example, the processing unit 106 may calculate the distance between diametrically opposed

points within the same element; or the distance between corresponding points in different elements (i.e. between E_i and E_{i+j}).

The processing unit 106 has been described with particular emphasis on the analysis of the element E_i in one image to calculate the metrics. It will be understood that the processing unit 106 may analyse multiple images. The different images may contain the same element E_i . The processing unit 106 may be configured to compare the scores generated for a particular element E_i in dependence on the analysis of the different images. If a discrepancy is detected between the scores, this can be flagged up as a potential problem with image quality or similar.

This may enable the quality of the different images to be checked. Similarly, particularly with fine-grained elements E_i , the scores calculated for elements E_i disposed proximal to each other are typically related. If an expected relationship is identified, this can be flagged as a potential image quality issue. These techniques may enable identification of an image of one or more of the element E_i , which is more likely to be correct.

The processing unit 106 may analyse the elements E_i visible in one or more images to infer the behaviour of elements E_i which are not visible. For example, a score may be estimated for an unsighted element E_i in dependence on a calculated score for at least one element E_i disposed adjacent to or proximal to the unsighted element E_i . Further analysis may be performed to build up a complete model of the left ventricle 202 based on the available views. The resulting model may enable scoring to be inferred from one or more nearby elements E_i .

The statistical analysis described herein was performed within the R statistical computing environment (v3.4.1), making use of the ggplot2, dplyr, ggbeeswarm, Momocs, pathmapping, and ggpubr packages. Due to the imbalance in the number of observations between groups in each comparison, multivariate, non-parametric hypothesis tests were employed to compare group means using the npmv package. A type I error rate (α) of 0.05 was used for all comparisons.

The processing unit 106 has been described herein as calculating metrics for elements E_i corresponding to the segments of a standard model of the left ventricle 102. It will be understood that the techniques described herein do not require that the elements E_i correspond to the segments. For example, the elements E_i may be smaller than the segments of the standard model. The elements E_i may correspond to sub-segments of the standard model. By reducing the size of the elements E_i , the scoring may provide a more precise indication of the location of abnormalities in the cardiac cyclic motion. For example, it is envisaged that the scoring may indicate the location of an abnormal function within one of the

segments of the standard model, for example highlighting a position near a boundary of the segment or in a central location.

As described herein, a reference data model is generated by analysing a reference data set

5 comprising raw two-dimensional echocardiography data. In the embodiment described above, the reference data set comprises historic data comprising end diastolic images and end systolic images for a group of patients. In a further development, the reference data set used to generate the reference data mode may be updated iteratively. For example, the analysis of new echocardiographs may be incorporated into the reference data set to increase the 10 available data population. Thus, the reference data model may continue to be refined as additional data becomes available. The iterative development of the reference data model may allow for pathological changes and patient evolution.

The analysis described herein is performed independently for each element E_i . However, it

15 will be appreciated that the analysis may be modified to consider the relationship between a plurality of elements E_i . For example, the analysis may simultaneously score the motion of first and second elements E_i which are disposed adjacent to each other or in opposition to each other, for example on opposing sides of the left ventricle 202.

20 The reference data model described herein may also be modified in dependence on outcome data available in respect of some or all of the reference data set. The term “outcome data” is used herein to refer to diagnostic information. The outcome data is associated with a corresponding record or set of data in the reference data set. The diagnostic information may, for example, relate to angiographic data and/or cardiac events for a patient. The outcome data 25 may indicate whether the patient had a positive or negative diagnosis for a cardiac condition, for example the presence or absence of coronary artery disease, during an elapsed time interval. The outcome data may, for example, be generated one (1) year, two (2) years or three (3) years after acquisition of the echocardiography data. A weighting of the data within the reference data set may be adjusted in dependence on the outcome data. For example, a 30 weighting applied to the data within the reference data set for which outcome data is available may be increased or decreased to change the statistical significance thereof. The weighting may be adjusted in dependence on the period of time elapsed between acquisition of the echocardiograph image and a subsequent diagnostic event. In a variant, the reference data model could be generated exclusively in dependence on data for which outcome data is 35 available. The reference data model could be generated exclusively in dependence on data for which the outcome data indicates the presence or absence of a particular condition, such

as coronary artery disease. The outcome data may be used to filter the reference data set to generate different reference data models.

The processing unit 106 may be configured also to provide a diagnostic function to generate

5 a diagnostic output. A diagnostic system is disclosed in the Applicant's International patent application PCT/GB2017/051720, the contents of which are incorporated herein in their entirety by reference. It has been recognised that the diagnostic function may utilise the outcome data described herein. The diagnostic function may also rely on one or more of the metrics generated for the wall motion score. By way of example, the rectangularity of each

10 element E_i may be used as a feature in the diagnostic model. The use of outcome data when generating a diagnostic model may help to take account of different disease proportions and characteristics over time and/or at different medical sites. For example, different sites may record different proportions of positive ('Disease') to negative ('Normal') outcomes. By utilising the outcome data in generating a diagnostic model, allowances may be made for these types

15 of variations. The results of stress echo test (as determined by a cardiologist during/shortly after the test) may not always be accurate. An analysis undertaken by the Applicant of one (1) year outcome accuracy has shown an average inaccuracy of 7.2% in stress echo results across multiple data sets. By referencing outcome data over a period of time, the accuracy of the diagnostic model may be improved, thereby enabling more accurate prediction of whether

20 or not an individual will go on to develop a disease, such as coronary artery disease. The use of outcome data is believed to be patentable independently. This enhanced diagnostic functionality will now be described as a development of the previous embodiment. Like reference numerals are used for like components.

25 As described herein, each end systole contour point 232-n is paired with a corresponding one of the end diastole contour points 242-n in the end diastole contour data set 244. The resulting pairs of end systole and end diastole contour points 232-n, 242-n represent changes in the motion of the wall of the heart 200 during a cardiac cycle. Once the end systole and end diastole contour points 232-n, 242-n have been identified, their x and y coordinates in the

30 Cartesian coordinate system may be stored in the memory 110, for example as an end systole coordinate set including the coordinates of the points on the end systole image and an end diastole coordinate set including the coordinates of the points on the end diastole image. The processor may be configured to calculate, from the two coordinate sets, the transformation in geometry of the left ventricle 202 between end systole and end diastole.

35 The processing unit 106 is configured to calculate values for various parameters that quantify the movement of the left ventricle 202 between end systole and end diastole. The calculation

may include working out how far each point has moved in each of the x and y directions, by working out the change in position (End diastole – End systole) along both the x axis and the y axis. This gives a set of x axis movements Δx and a set of y axis movements Δy for each corresponding pair of end systole and end diastole contour points 232-n, 242-n. Each of these

5 values may be a simple distance with no indication of direction. The mean change of all the points in both the x axis (Δx) and y axis (Δy) may then be calculated separately so as to provide an average Δx value or x direction movement ΔX , and an average Δy value or y direction movement ΔY for the entire left ventricle 202. If each of the individual movement values are purely distance, without any indication of whether they are in the positive or 10 negative x or y direction, then these averages will describe the total amount of movement, but not give an indication of the direction or of whether different parts of the LV wall are moving in the same direction or opposite directions.

Another parameter that may be calculated for each pair of end systole and end diastole 15 contour points 232-n 242-n is the mean of the x and y direction movements Δx and Δy , where the mean value for each point $\Delta xy = (\Delta x + \Delta y)/2$. The mean of all the values of Δxy for all points can then be calculated to a value for the entire ventricle ΔXY . This calculation is similar to the calculation of shear strain and is therefore referred to herein as the shear transformation. It will be appreciated that, for a given distance of movement, this parameter will be largest for 20 movements at 45 degrees to both of the x and y axes, and smallest for movements along one of the axes.

A further parameter that can be calculated is similar to the principal transformation that can be calculated from x and y strain components, and is therefore referred to herein as the 25 principal transformation, given by

$$\text{Principal transformation} = C1(\Delta X + \Delta Y - \sqrt{(\Delta X + \Delta Y)^2 + C2\Delta XY^2})$$

where C1 and C2 are constants. The constant C1 may, for example, be $\frac{1}{2}$ and the constant C2 may be 4. These values are used in the examples described below.

30 This transformation is closely related to the shear transformation and therefore tends to vary in a similar way to that parameter, but has a negative value indicating contraction of the heart. However, as indicated by the test results below, the principal transformation value can give a more reliable diagnosis in some cases, in particular of coronary artery disease (CAD).

35 It will be appreciated that each of these parameters relates to changes between end systole and end diastole in a single coronary cycle. However in stress echocardiography, (or

corresponding tests carried out with other imaging methods) there will be one value for each parameter for the heart at rest and one value for the heart at stress. Comparing those values, for example determining the difference between them, gives further information about heart function that can be used in diagnosis.

5

Once the x and y movements, and shear and principal transformation values have been calculated, the processor is then configured to compare these with reference values stored in the memory 110 to make a diagnosis of one or more specific heart conditions, and to generate a diagnostic output. The output may be a simple binary output indicating a positive or negative diagnosis. The processor unit 106 may be arranged to display the output on the display 112. Alternatively, or in addition, it may be arranged to store the output as data in association with the images on which it was based, for example by adding output data, indicative of the diagnosis, to a file in which the images are stored.

10

The reference values may be determined by means of a learning algorithm which, for example, can be run on the processor unit 106, and which uses a database of stress echo images with associated diagnoses as determined by conventional methods, which may be stored in the memory 110. Specifically, the database may include a large number of sets of images, each set comprising an end systole image and an end diastole image for both rest condition and stress condition, together with, for each set of images, an associated diagnosis, such as a positive or negative diagnosis for coronary artery disease (CAD). The learning algorithm may be arranged to analyse the images to calculate values of the various parameters described above, and then to determine the correlation between the diagnosis and the values of each of the various parameters.

15

Analysis was carried out on sample images from seventy (70) subjects. All results generated were from an apical four chamber view. Firstly the values were compared for positive and negative outcomes as determined from the DSE results. Then the comparison was repeated with the DSE results corrected for confirmed false positives in the DSE results.

20

Table 1 Shows values of the principal transformation (in mm), shear transformation value (in mm), and mean ΔX (in mm) at rest and stress for DSE outcome (1= Pos, 2= Neg) in the Apical four Chamber view.

5

Group Statistics

DSE_Result	N	Mean	Std. Deviation	Std. Error Mean
Stress_Prin	1.00	9	-6.8214	4.08788
	2.00	61	-8.9260	2.20018
Rest_Prin	1.00	9	-7.7332	3.86497
	2.00	61	-9.3163	2.41589
Rest_Shr	1.00	9	17.7267	9.16943
	2.00	61	21.5356	5.50610
Stress_Shr	1.00	9	17.0074	8.06969
	2.00	61	22.2608	4.56871
Rest_X	1.00	9	18.8694	11.02116
	2.00	61	21.8492	6.65078
Stress_X	1.00	9	19.9334	9.80639
	2.00	61	25.8710	7.43965

Table 2 Shows means of Principal transformation value (in mm), Shear transformation (in mm) and X transformation (in mm) at rest and stress for Adjusted DSE outcome (1= Pos, 2= Neg).

10

Group Statistics

Adjusted_DSE	N	Mean	Std. Deviation	Std. Error Mean
Stress_Prin	1.00	7	-4.4716	1.29120
	2.00	63	-9.1203	2.24588
Rest_Prin	1.00	7	-5.3352	1.21275
	2.00	63	-9.5325	2.44136
Rest_Shr	1.00	7	12.0645	2.74525
	2.00	63	22.0438	5.58342
Stress_Shr	1.00	7	12.2348	3.81629
	2.00	63	22.6243	4.44025
Rest_X	1.00	7	11.6937	2.73459
	2.00	63	22.5519	6.84823
Stress_X	1.00	7	14.1727	4.81157
	2.00	63	26.3226	7.29318

Table 3 shows independent samples T-Test for variables vs adjusted DSE.

		Independent Samples Test				
		Levene's Test for Equality of Variances				
		F	Sig.	t	df	Sig. (2-tailed)
Stress_Prin	Equal variances assumed	1.705	.196	5.356	68	.000
	Equal variances not assumed			8.240	10.596	.000
Rest_Prin	Equal variances assumed	2.355	.130	4.466	68	.000
	Equal variances not assumed			7.604	12.377	.000
Rest_Shr	Equal variances assumed	2.106	.151	-4.644	68	.000
	Equal variances not assumed			-7.961	12.527	.000
Stress_Shr	Equal variances assumed	.194	.661	-5.942	68	.000
	Equal variances not assumed			-6.715	7.923	.000
Rest_X	Equal variances assumed	5.696	.020	-4.136	68	.000
	Equal variances not assumed			-8.065	16.500	.000
Stress_X	Equal variances assumed	.927	.339	-4.290	68	.000
	Equal variances not assumed			-5.963	9.395	.000

From the values of the various parameters obtained from the sample data, machine learning may be used to determine the accuracy of each parameter as an indicator of adjusted

5 Dobutamine stress echo (DSE) outcome. Using the data above, a J48 pruned decision tree with 10 fold cross validation method was used to classify the data. The accuracy of each parameter as an indicator of diagnostic outcome is summarized in the tables below, in which the following abbreviations are used: TP = true positive; FP = false positive; FN = false negative; TN = true negative; PPV = positive predictive value; and NPV = negative predictive value.

Table 4 Accuracy of Consultant Interpretation

J48 Accuracy = 94.3%	TP = 6	FN = 1
	FP = 3	TN = 60
Sensitivity = 85.7%	PPV = 66.7%	
Specificity = 95%	NPV = 98.4%	

Table 5 Accuracy of Stress Principal Transformation for Adjusted DSE outcome

J48 Value = -5.95	TP = 7	FN = 0
Accuracy = 95.7%	FP = 3	TN = 60
Sensitivity = 100%	PPV = 70%	
Specificity = 95.2%	NPV = 100%	

Table 6 Accuracy of Rest Principal Transformation for Adjusted DSE outcome

J48 Value = -6.92	TP = 5	FN = 2
Accuracy = 88.6%	FP = 6	TN = 57
Sensitivity = 71.4	PPV = 45.5 %	
Specificity = 90.5%	NPV = 96.6%	

5 **Table 7** Accuracy of Stress Shear Transformation for Adjusted DSE outcome

J48 Value = 15.85	TP = 6	FN = 1
Accuracy = 95.7%	FP = 2	TN = 61
Sensitivity = 85.7%	PPV = 85.7	
Specificity = 96.8 %	NPV = 98.4	

Table 8 Accuracy of Rest Shear Transformation for Adjusted DSE Outcome

J48 Value = 15.35	TP = 5	FN = 2
Accuracy = 91.4%	FP = 4	TN = 59
Sensitivity = 71.4	PPV = 55.6%	
Specificity = 93.7 %	NPV = 96.7%	

Then from all of the variables, using machine learning, a decision tree which is shown in Figure 10 13 was derived to provide accurate diagnosis from the data. The decision tree defines a series of decision points, each of which defines a reference or threshold value of a parameter. The decision tree outlines a simple algorithm which operates as follows. Firstly the principal transformation of the left ventricle 202 as described above is determined for the stress condition of the heart. If the transformation is less than -5.95mm (i.e. a negative value with 15 magnitude greater than 5.95mm) then the diagnosis is negative. If the value is greater than -5.95mm (i.e. a negative value with magnitude greater than 5.95mm) then difference in principal transformation between rest and stress conditions is greater than 12.278053mm then the diagnosis is negative, but if it is less than that distance, the diagnosis is positive. It will be appreciated that the structure of the decision tree, and the reference or threshold values at 20 each decision point in the decision tree, will depend on the diagnosis that is to be made.

The processing unit 106 described above implements a fixed (static) diagnostic model for diagnosing coronary artery disease. As illustrated in Figure 13, the decision tree defines a series of decision points, each of which defines a reference or threshold value of a parameter. The processing unit 106 may implement a dynamic diagnostic model. The reference or 5 threshold values at each decision point in the decision tree may be modified dynamically, for example to reflect the new echocardiographic data and/or outcome data. At least in certain embodiments, this may provide improved diagnostic functions.

10 The new echocardiographic data may be incrementally added to the existing set of reference data. The new data is used to expand the data population and may progressively change the diagnostic model. The reference or threshold values used in the decision tree may be updated to reflect the available echocardiographic data. The iterative development of the reference data allows the diagnostic model to change with respect to time. It will be understood that the 15 decision tree described herein may be replaced with other analysis tools, such as a supervised machine learning model.

20 The outcome data comprises diagnostic information for each patient, for example relating to angiographic data and/or cardiac events. The outcome data in the present embodiment indicates whether the presence or absence of coronary artery disease was detected during an elapsed time interval after acquisition of the end systole image and end diastole image used in the reference data set. The outcome data may, for example, be generated one (1) year, two (2) years, three (3) years or longer after acquisition of the echocardiography data. The outcome data in the present embodiment is generated one (1) year after acquisition of 25 the echocardiography data. The outcome data is compiled by considering any angiographic data and cardiac events that have taken place during the elapsed time interval. It will be understood that the outcome data continues to evolve with respect to time. The outcome data may, therefore, be updated on an ongoing basis, for example on an annual basis or when a classification changes. By updating the outcome data, the diagnostic tools and diagnostic models generated in dependence on the reference data may be adjusted dynamically to 30 represent pathological changes and patient evolution.

35 In order to implement the dynamic diagnostic model, a classification model is built using a supervised machine learning algorithm. The outcome data is used to label the reference data accessed by the machine learning algorithm. The machine learning algorithm uses the labels to distinguish between different classifications. In the present embodiment, the classifications correspond to the presence or absence of coronary artery disease. Alternatively, or in addition, the classifications may grade a particular condition, for example in dependence on an

identified stenosis level or percentage. It will be understood that the classifications may distinguish between other conditions. When generating the diagnostic models, the machine learning algorithm may adjust the relative weighting of the reference data in dependence on the labels derived from the outcome data. At least in certain embodiments, updating the 5 reference data in dependence on the outcome data may provide improved diagnostic accuracy based on the stress echocardiograms.

In order to build a diagnostic model, a set of features are calculated from the contour data. The features are calculated per-segment (for example by analysing one or more of the 10 elements E_i described herein) and optionally in respect of the entire left ventricle 202. The available feature-set is analysed to identify those features that are most relevant. The most pertinent features may thereby be identified to build the diagnostic model. In the case of a random forest (which consists of multiple decision trees), the identified features form the decision nodes. The most relevant features may vary across geographic regions and/or 15 change as the disease evolves, the features identified for use in the model may change. Even if the features remain the same, the thresholds and weightings may change. As shown in Tables 1 and 2 herein, the top feature remains unchanged as the ejection fraction at peak stress. However, the next most important features changes for the different conditions. In the first data set (Table 1), the volume change between end-systole and end-diastole is the next 20 most relevant. However, in the combined dataset, the area of a specific segment at rest in the two-chamber view is the next most relevant. In order to train the model, the reference data needs to be labelled. In view of the potential inaccuracies, using the results of a stress echo (as determined by a cardiologist) as the label will not necessarily lead to an accurate model. The use of outcome data that is collected a period of time after the acquisition of the reference 25 data (for example, one (1) year after acquisition of the echocardiograph images), at least some of these deficiencies can be overcome or ameliorated.

The outcome data can be collected for different periods of time. The outcome data can, at least in certain embodiments, provide an indication of how far in advance the effects of 30 coronary artery disease can be identified. Moreover, multiple classes of labels can be used to predict different disease severity. As more outcome data is accumulated, the diagnostic model is updated to help ensure that the classification remains as accurate as possible due to the possibility of disease evolution and population changes. This can be done by retraining the entire model every time new outcome data is received. In practice, this may prove time-consuming. As an alternative, incremental machine learning techniques can be implemented 35 by the processor to continually update the diagnostic model.

The implementation of the classification model will now be described with reference to a first reference data set and a second reference data set. The first data set comprises a first set of one hundred and twenty-four (124) stress echocardiograms (collected in Oxford between May 2011 and August 2013). The second data set comprises a set of three hundred and thirty-nine

5 (339) stress echocardiograms from a separate study (collected between March 2015 and August 2016 in six (6) different hospitals across the Thames Valley). The outcome data is compiled one (1) year after acquisition of the stress echocardiograms. The outcome data generates a binary outcome value. In particular, an outcome is considered positive if during the elapsed one (1) year interval one of the following events is identified:

10 (i) a cardiac event (e.g. myocardial infarction);
 (ii) an angiogram which showed greater than 70% stenosis.

The outcome is considered negative if neither of the aforementioned events (i) or (ii) occurred in the elapsed one (1) year interval. In the first data set, ten (10) positive outcomes were identified, and in the second data set thirteen (13) positive outcomes were identified.

15 The Boruta package from the R statistical computing environment to assess the most relevant features for predicting an outcome. The Boruta package performs feature selection by comparing the importance of attributes to those possible at random. A standard implementation comprising a random forest with 500 trees was implemented. Table 9 details
 20 the most important features and their mean importance score for the first dataset. The second data set was added to the first data set. Table 10 details the most important features and their mean importance score for the combined first and second data sets. The most relevant features change as more data is available for processing. This demonstrates that the classification model may change with the addition of more reference data. It is believed that
 25 these changes would be more pronounced if the additional reference data is acquired at a later date and/or over a more widespread geographical area. Although the use of a random forest model has been described herein, it will be understood that another model could be used, or indeed an ensemble of models.

30 **Table 9** Most relevant features using the first data set

Feature	Mean importance
EF_P Ejection fraction at peak stress	6.23
ES_P_to_ED_P Ratio of end-systolic to end-diastolic peak volume	6.20
rect_segment_4_R_2C Rectangularity of the apical anterior segment at rest	5.42

solid_segment_4_R_2C Solidity of the apical anterior segment at rest	5.41
norm_area_segment_4_R_2C Normalised area of the apical anterior segment at rest	5.35
ES_P_to_ED_P_2C Ratio of the end-systolic to end-diastolic 2 chamber area at peak stress	5.00
ES_P_to_ED_P_4C Ratio of the end-systolic to end-diastolic 4 chamber area at peak stress	4.70
P_ES End-systolic volume at peak stress	4.53
total_ES_area_P_2C 2 chamber end-systolic area at peak stress	3.86
dy_8_P_4C Euclidean distance of the eighth point in 4 chamber at peak stress	3.75

Table 10 Most relevant features using combined data from the first and second data sets

Feature	Mean importance
EF_P Ejection fraction at peak stress	6.59
norm_area_segment_4_R_2C Normalised area of the apical anterior segment at rest	6.53
ES_P_to_ED_P_4C Ratio of the end-systolic to end-diastolic 4 chamber area at peak stress	6.38
norm_area_segment_4_P_4C Normalised area of the apical anterior segment at rest	5.05
total_ES_area_P_4C 4 chamber end-systolic area at peak stress	4.30
ES_P_to_ED_P Ratio of the end-systolic to end-diastolic area at peak stress	4.10
prin_trans_P_4C Principal strain in the 4 chamber view at peak stress	3.96
solid_segment_4_R_2C Solidity of the apical anterior segment at rest	3.96
norm_d_segment_6_P_4C Normalised average distance in the basal lateral segment	3.94
ES_P_to_ED_R_4C Ratio of the end-systolic to end-diastolic 4 chamber area at peak stress	3.91

The implementation of a continued learning strategy capable of incorporating new reference data may provide a more robust and accurate diagnostic model may be achieved. By incorporating the new reference data incrementally, the need to retrain the entire model may be reduced or avoided each time new data becomes available (which can prove a time-

consuming process, particularly as the size of the reference data set increases). Moreover, the diagnostic model can adapt to changing disease characteristics over time. This is particularly important as the most relevant biomarkers may change over time due to the changing environments and lifestyles of the population, and the model needs to adapt to 5 account for these. The dynamic diagnostic model can adapt to changing facets and characteristics of cardiovascular disease, thereby providing a robust and accurate prediction model.

10 The dynamic diagnostic model described herein utilises outcome data acquired over a one (1) year period. It will be understood that the outcome data may be accumulated over different periods of time. By combining the outcome data over a longer time period, the predictive power of the dynamic diagnostic model over a longer time period may be assessed.

15 The present application has been described with reference to cardiovascular disease. However, it will be understood that the methods and apparatus described herein may have other applications. For example, diagnostic tools may be developed to adapt to the changing imaging biomarkers for a tumour if the environment changes and the tumour size, appearance or calcification changes. Furthermore, the techniques may be applicable in imaging systems other than echocardiographs.

20

It will be appreciated that various modifications may be made to the embodiment(s) described herein without departing from the scope of the appended claims.

CLAIMS:

1. A system for generating a diagnostic model, the system comprising a processor configured to:
 - 5 analyse a plurality of reference data sets, each reference data set comprising at least one image, the analysis comprising identifying at least one feature in each image;
 - calculate at least one metric in dependence on the at least one identified feature;
 - acquire outcome data associated with at least some of the reference data sets; and
 - compile the diagnostic model in dependence on the at least one calculated metric
- 10 and the associated outcome data.
2. A system as claimed in claim 1, wherein each reference data set comprises first and second images, the processor being configured to analyse each reference data set to:
 - identify at least one first feature in the first image, and identify at least one second feature in the second image, each at least one first feature being paired with a corresponding one of the at least one second feature; and
 - 15 compare each pair of corresponding first and second features to identify one or more difference therebetween.
- 20 3. A system as claimed in claim 2, wherein the processor is configured to calculate the at least one metric in dependence on the one or more difference identified between each pair of corresponding first and second features.
4. A system as claimed in claim 2 or claim 3, wherein the first image comprises a first 25 end systolic image and the second image comprises a second end diastolic image, the one or more difference identified between each pair of corresponding first and second features represents a cardiac cyclic change.
5. A system as claimed in any one of the preceding claims, wherein the processor is 30 configured to label the images in dependence on the outcome data.
6. A system as claimed in claim 5, wherein the label distinguishes between different classifications.
- 35 7. A system as claimed in any one of the preceding claims, wherein the outcome data comprises diagnostic information.

8. A system as claimed in claim 7, wherein the diagnostic information comprises a record of a cardiac event, such as a myocardial infarction.

9. A system as claimed in claim 7 or claim 8, wherein the outcome data comprises a record of stenosis greater than a threshold value.

10. A system as claimed in any one of the preceding claims, wherein the processor is configured to update the diagnostic model when the outcome data becomes available or when the outcome data is updated.

10

11. A system as claimed in any one of the preceding claims, wherein the processor is configured to determine a weighting for at least some of the calculated metrics, each weighting being determined in dependence on the outcome data associated with a given one of the reference data sets.

15

12. A system as claimed in claim 11, wherein the diagnostic model is compiled in dependence on the calculated metrics and the associated weightings.

20

13. A system as claimed in any one of the preceding claims, wherein the processor is configured to implement a machine learning algorithm to generate the diagnostic model.

14. A system as claimed in any one of the preceding claims, wherein the processor is configured to analyse further reference data sets; and to update the diagnostic model in dependence on the analysis of the further reference data sets.

25

15. A method of generating a diagnostic model, the method comprising:
 analysing a plurality of reference data sets, each reference data set comprising at least one image, the analysis comprising identifying at least one feature in each image;
 calculating at least one metric in dependence on the at least one identified feature;
30 acquiring outcome data associated with at least some of the reference data sets; and
 compiling the diagnostic model in dependence on the calculated metrics and the associated outcome data.

30

16. A method as claimed in claim 15, wherein each reference data set comprises first and second images, the method comprising analysing each reference data set to:

identify at least one first feature in the first image, and identify at least one second feature in the second image, each at least one first feature being paired with a corresponding one of the at least one second feature; and

5 comparing each pair of corresponding first and second features to identify one or more difference therebetween.

17. A method as claimed in claim 16 comprising calculating the at least one metric in dependence on the one or more difference identified between each pair of corresponding first and second features.

10

18. A method as claimed in claim 16 or claim 17, wherein the first image comprises a first end systolic image and the second image comprises a second end diastolic image, the one or more difference identified between each pair of corresponding first and second features represents a cardiac cyclic change.

15

19. A method as claimed in any one of claims 15 to 18, wherein the outcome data is used to label the corresponding reference data sets.

20
20

20. A method as claimed in claim 19, wherein the label distinguishes between different classifications.

21. A method as claimed in any one of claims 15 to 20, wherein the outcome data comprises diagnostic information.

25

22. A method as claimed in claim 21, wherein the diagnostic information comprises a record of a cardiac event, such as a myocardial infarction.

23. A method as claimed in claim 21 or claim 22, wherein the outcome data comprises a record of stenosis greater than a threshold value.

30

24. A method as claimed in any one of claims 15 to 23 comprising updating the diagnostic model when the outcome data is updated and/or when new outcome data becomes available.

35

25. A method as claimed in any one of claims 15 to 24 comprising determining a weighting for at least some of the calculated metrics, each weighting being determined in dependence on the outcome data associated with a given one of the reference data sets.

26. A method as claimed in claim 25, wherein the diagnostic model is compiled in dependence on the calculated metrics and the associated weightings.

27 A method as claimed in any one of claims 15 to 26, wherein a machine learning
5 algorithm is implemented to generate the diagnostic model.

28. A method as claimed in any one of claims 15 to 27 comprising adding further
reference data sets incrementally, the method comprising analysing the further reference data
sets; and updating the diagnostic model in dependence on the analysis of the further reference
10 data sets.

29. A non-transitory computer-readable medium having a set of instructions stored
therein which, when executed, cause a processor to perform the method claimed in any one
of claims 15 to 28.

15

1/14

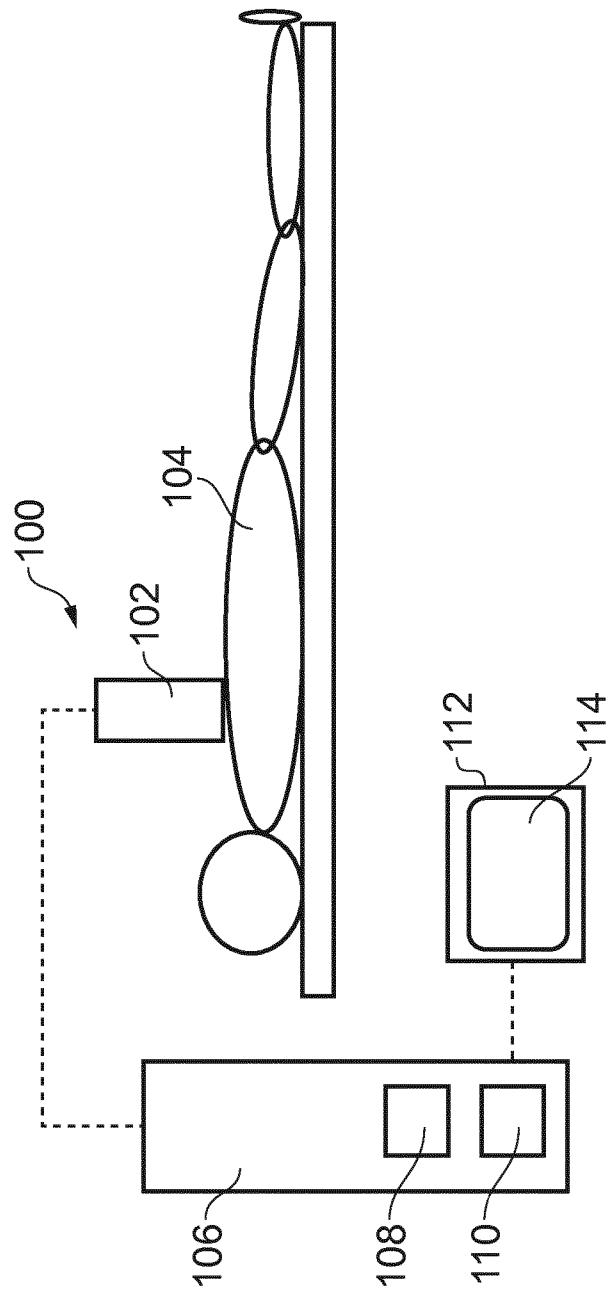


FIG. 1

2/14

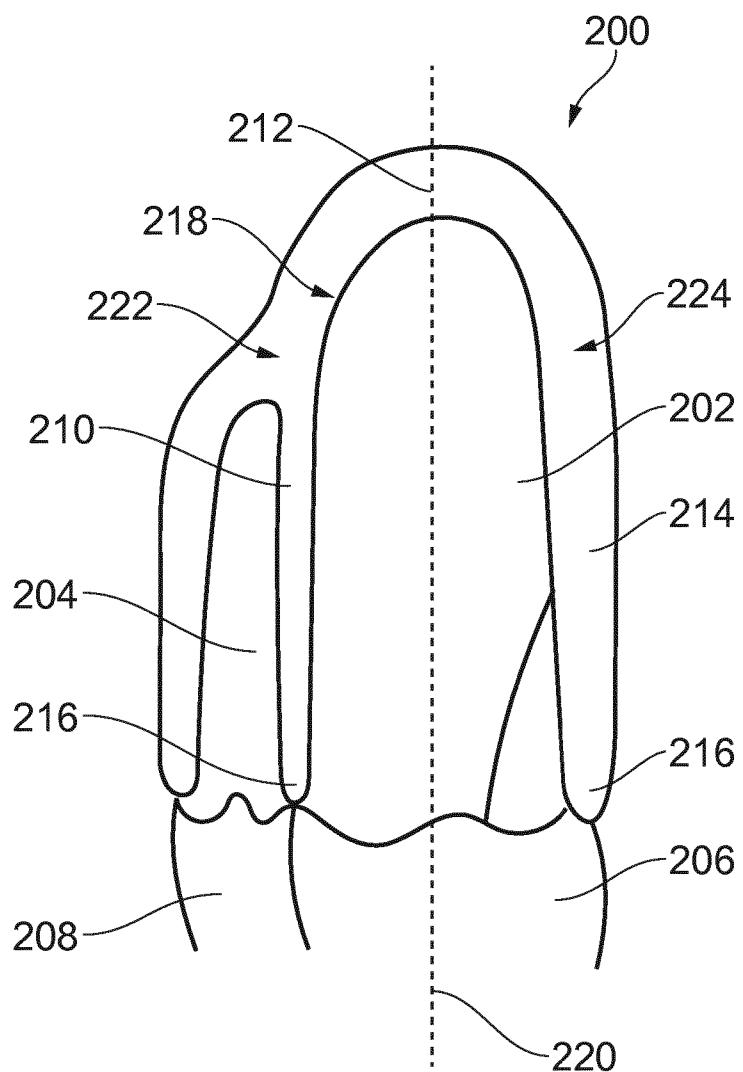


FIG. 2

3/14

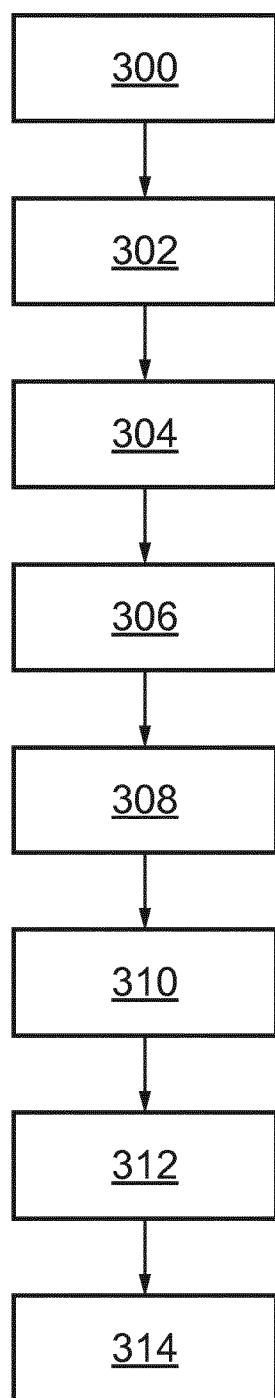


FIG. 3

4/14



FIG. 4A

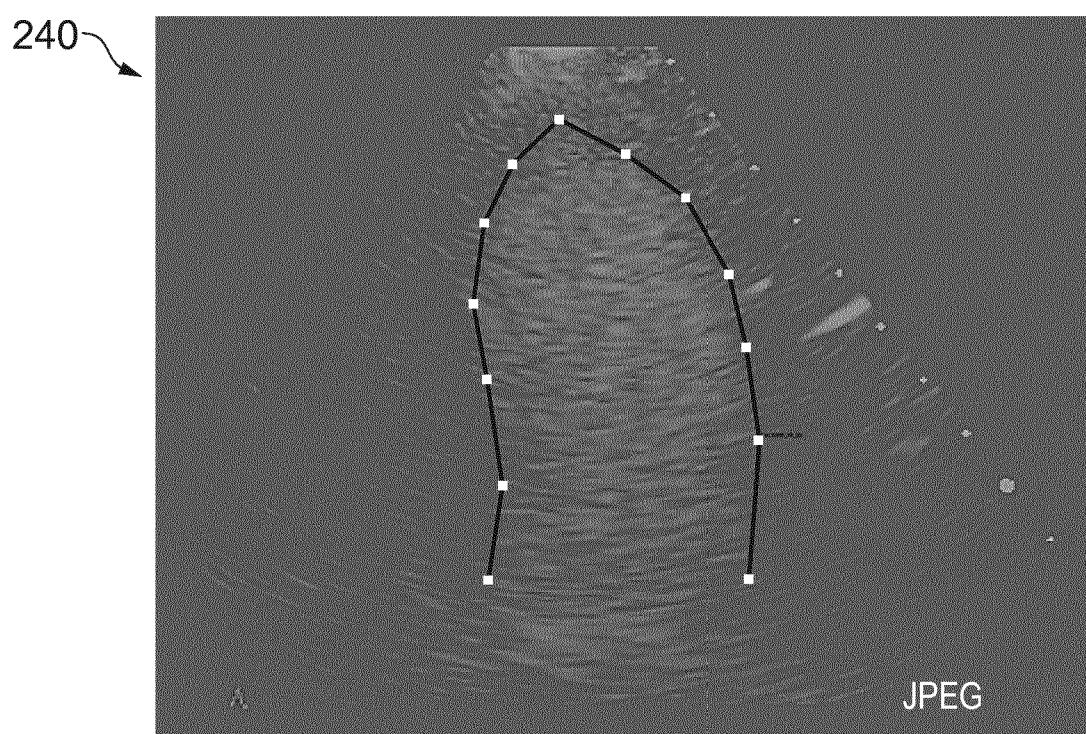


FIG. 4B

5/14

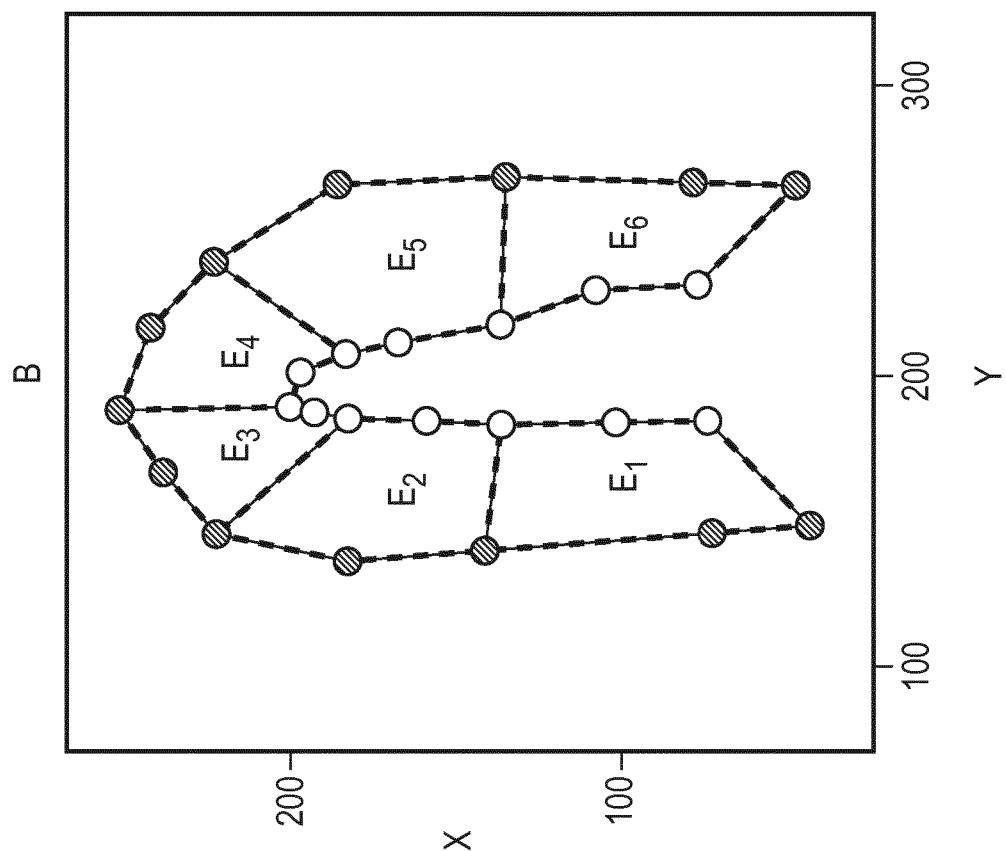


FIG. 5B

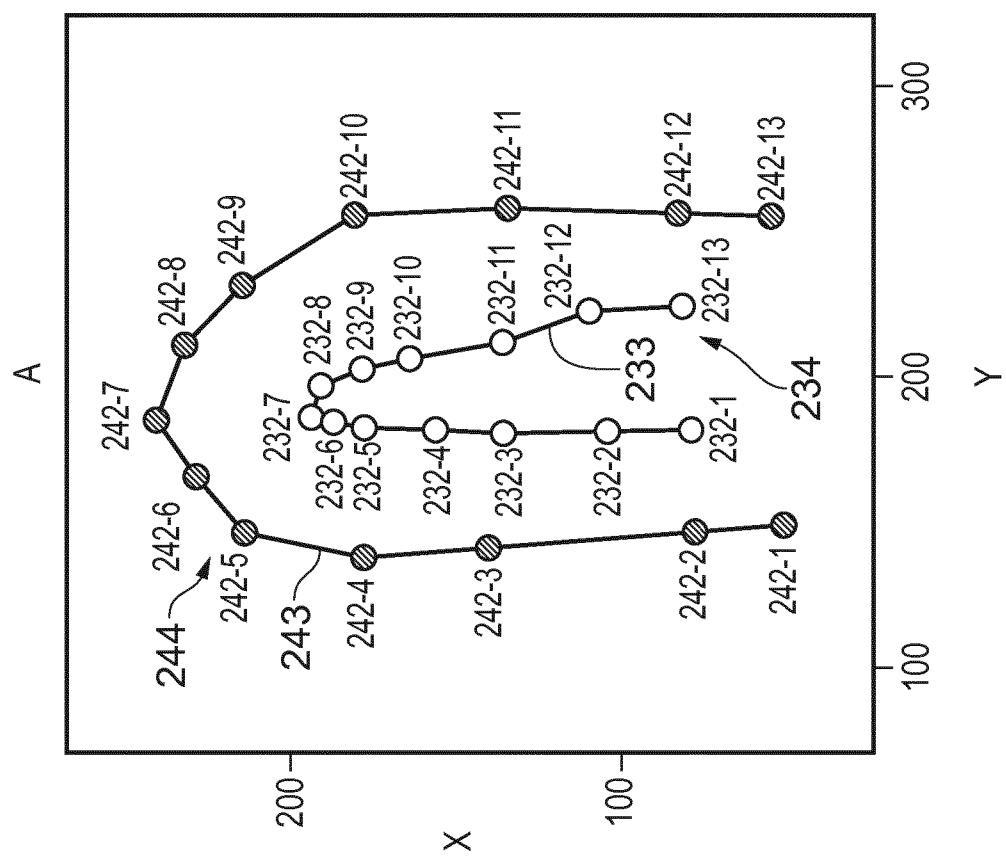


FIG. 5A

6/14

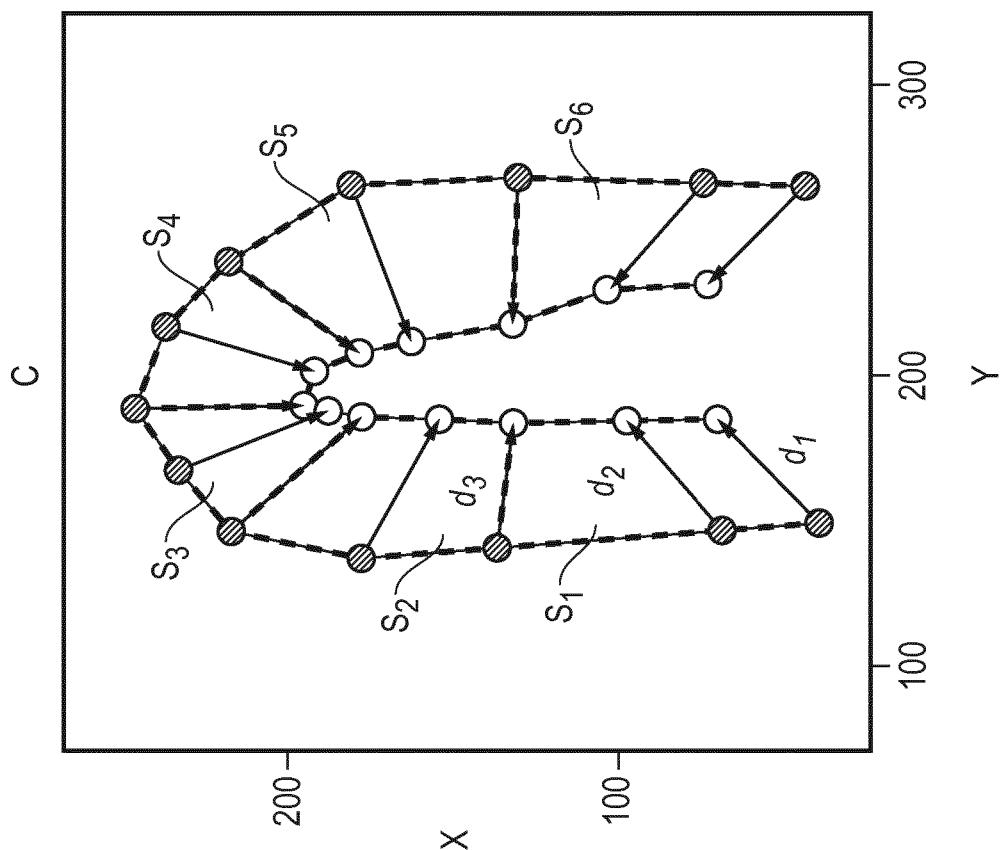


FIG. 6B

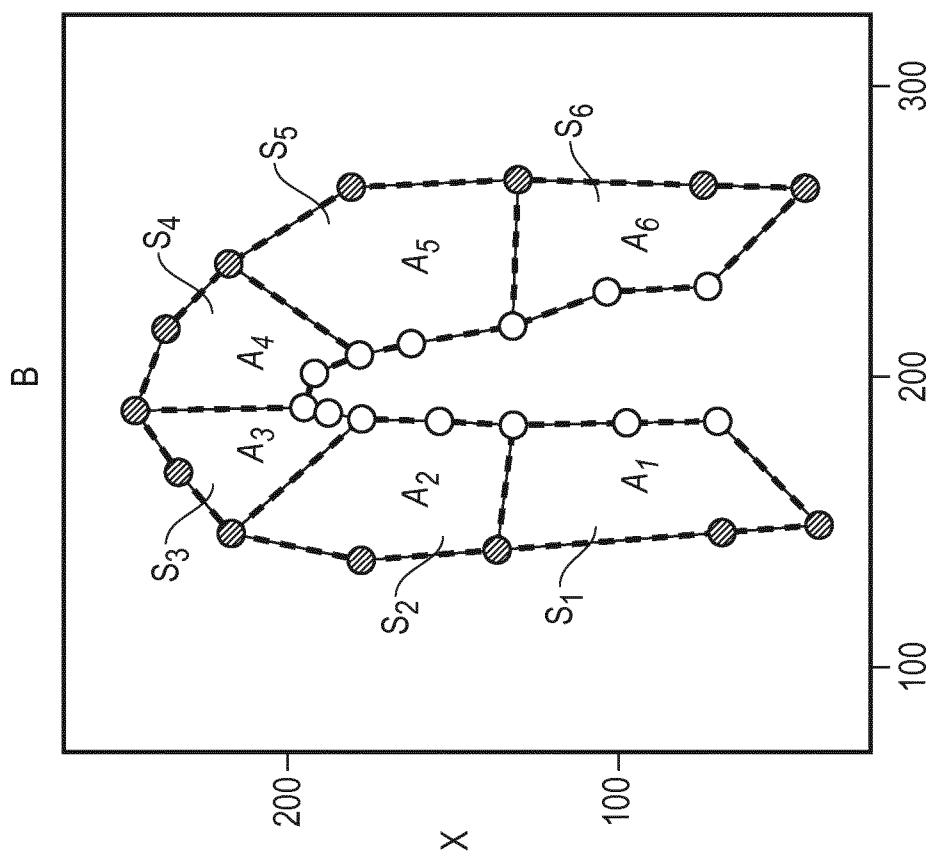


FIG. 6A

7/14

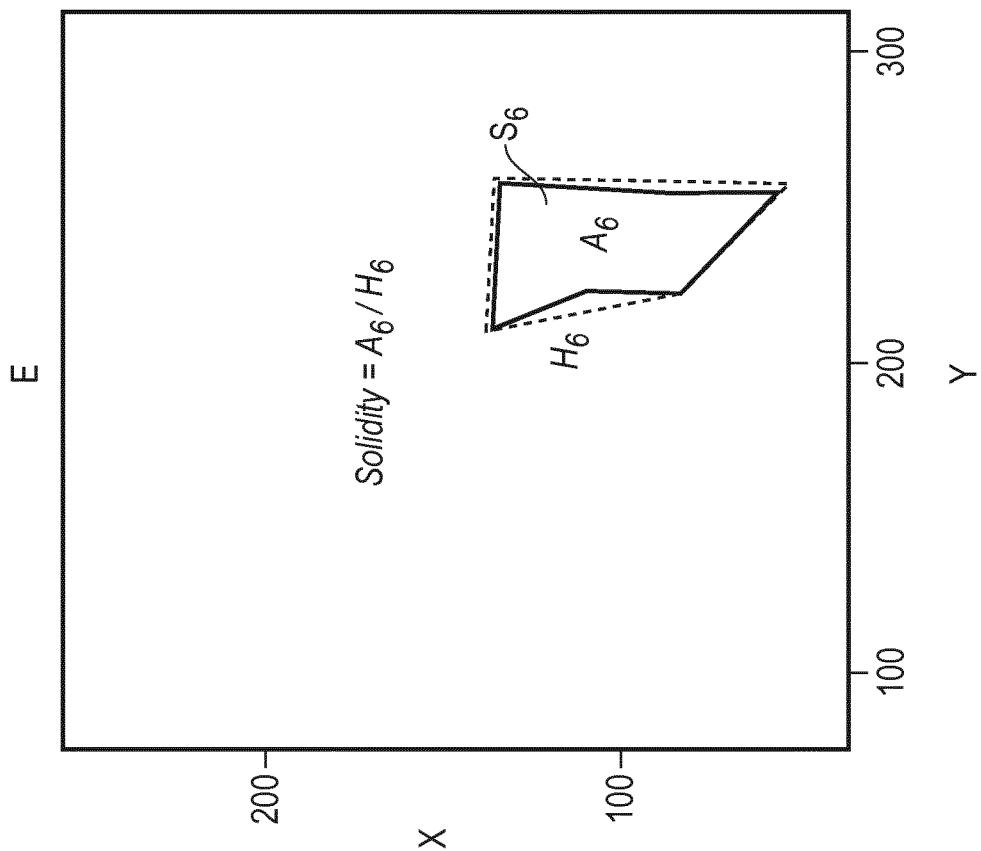


FIG. 6D

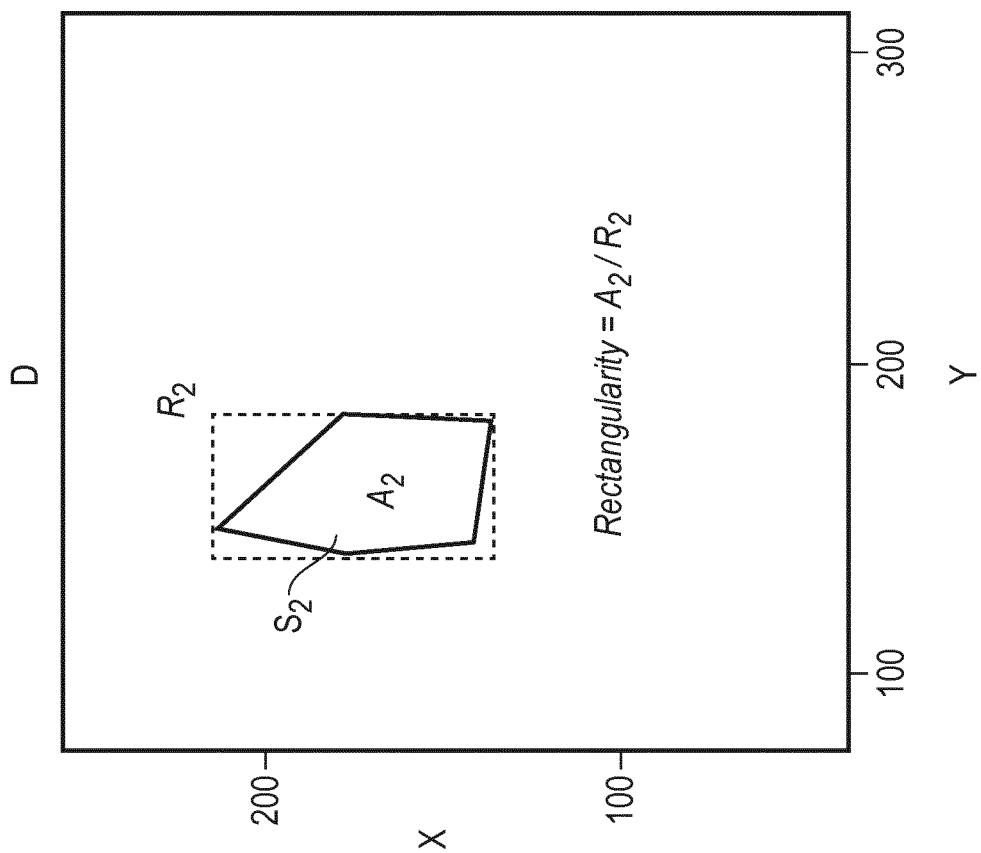


FIG. 6C

8/14

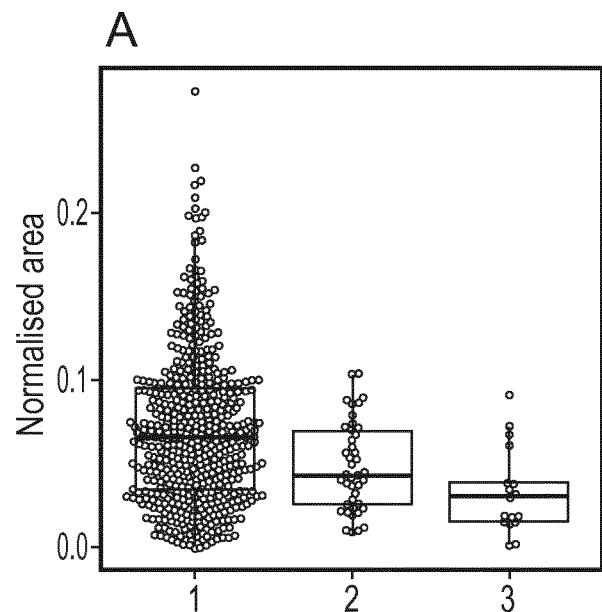


FIG. 7A

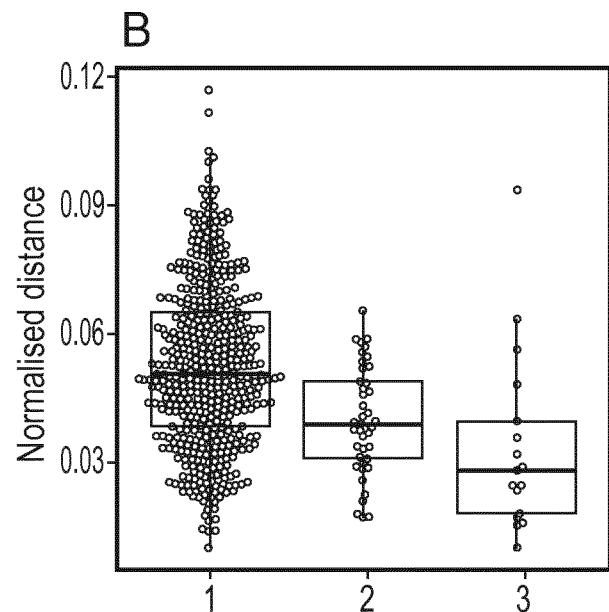


FIG. 7B

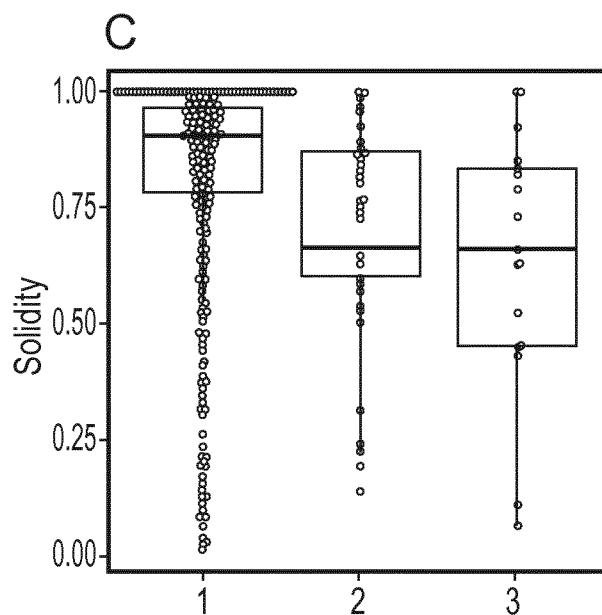


FIG. 7C

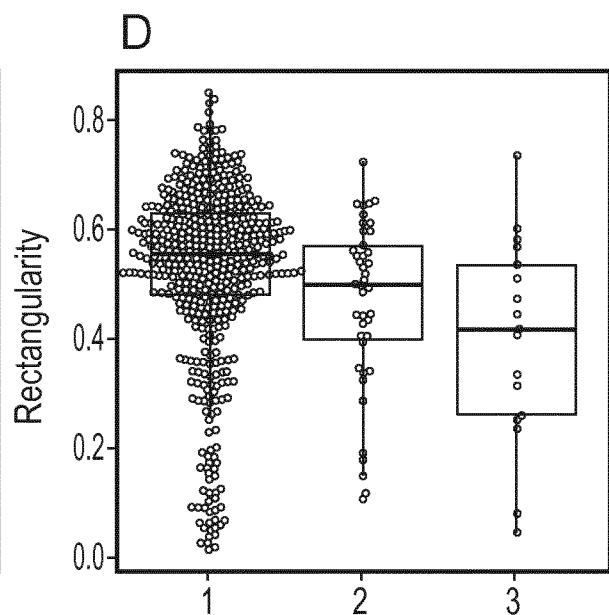


FIG. 7D

Wall motion score

9/14

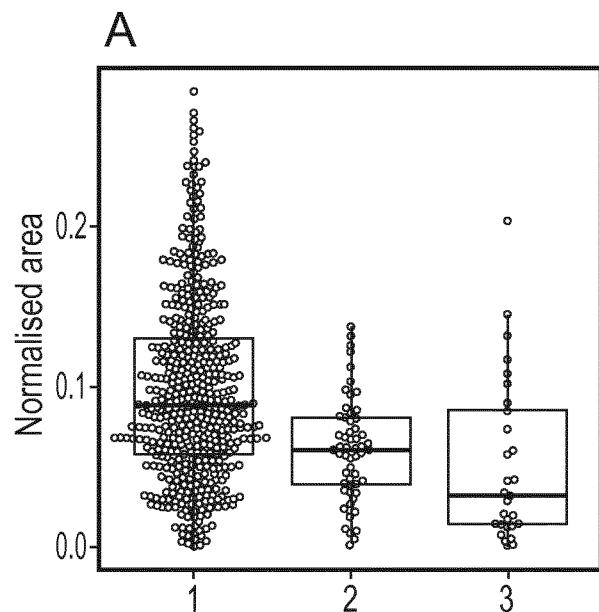


FIG. 8A

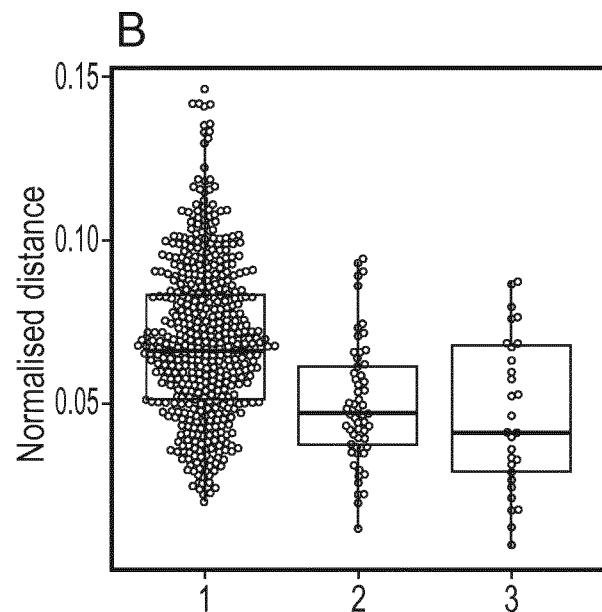


FIG. 8B

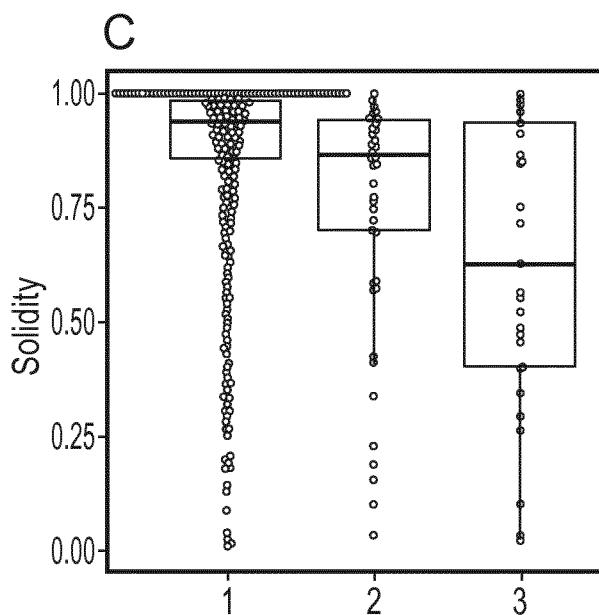


FIG. 8C

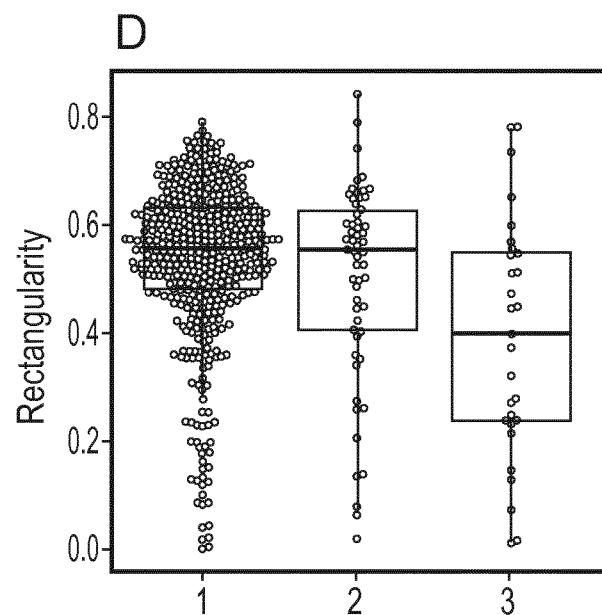


FIG. 8D

Wall motion score

10/14

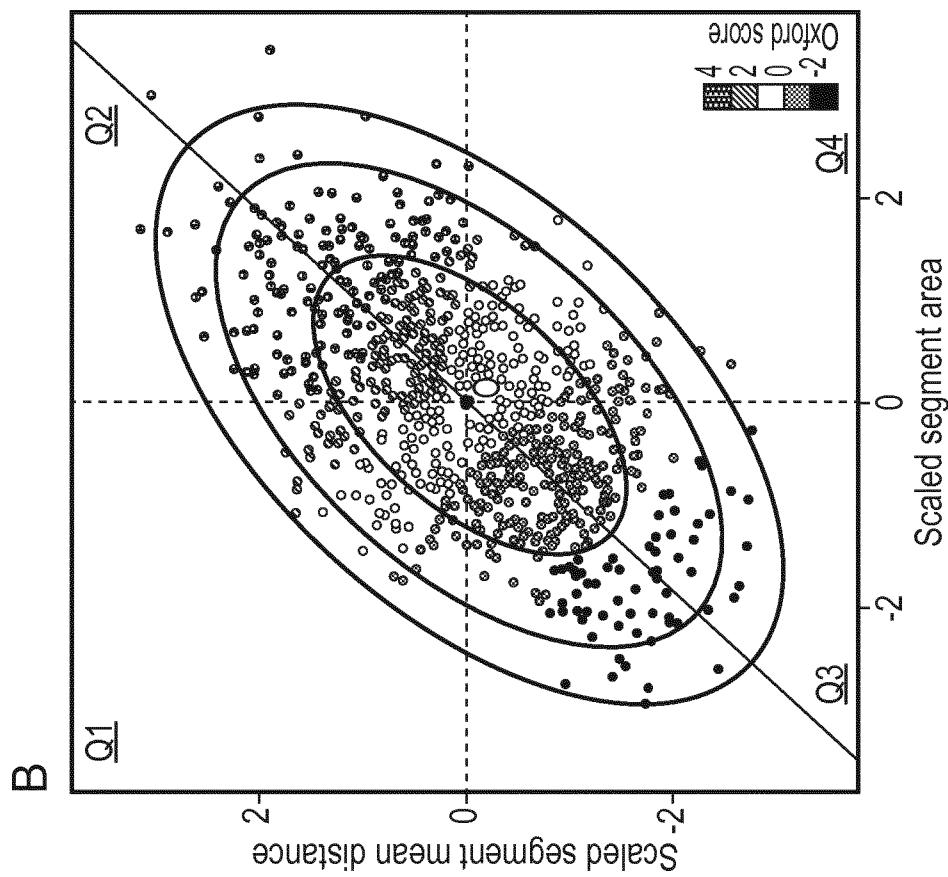


FIG. 9B

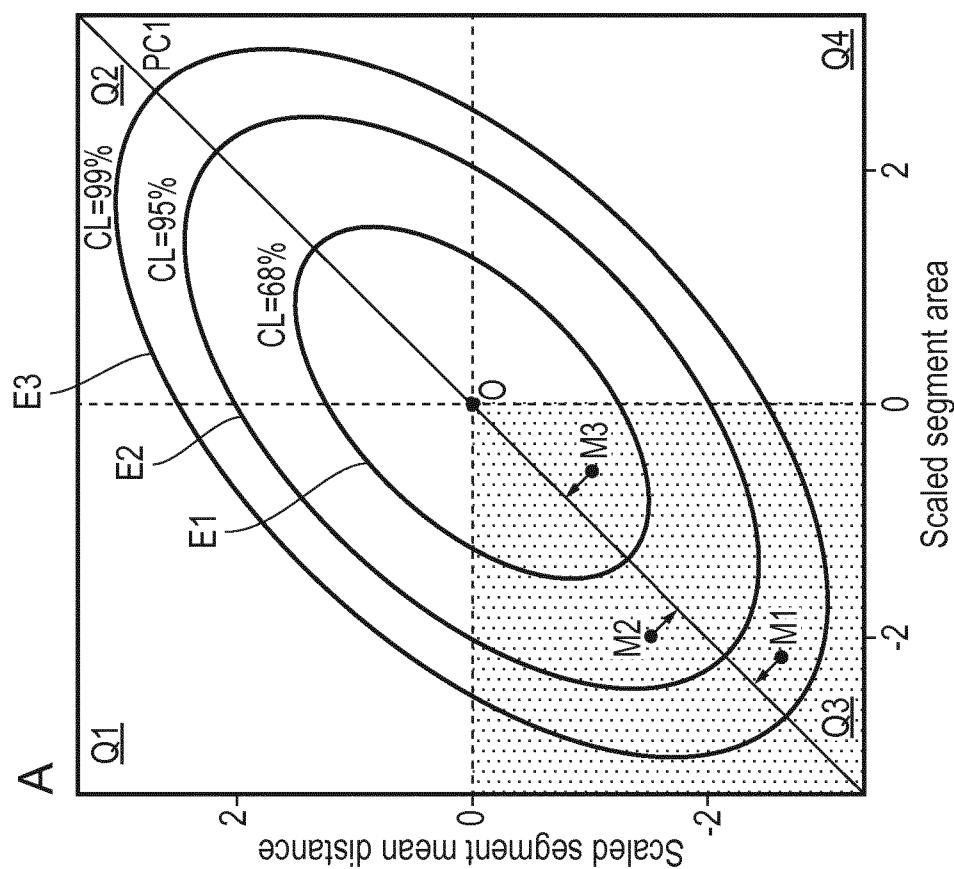


FIG. 9A

11/14

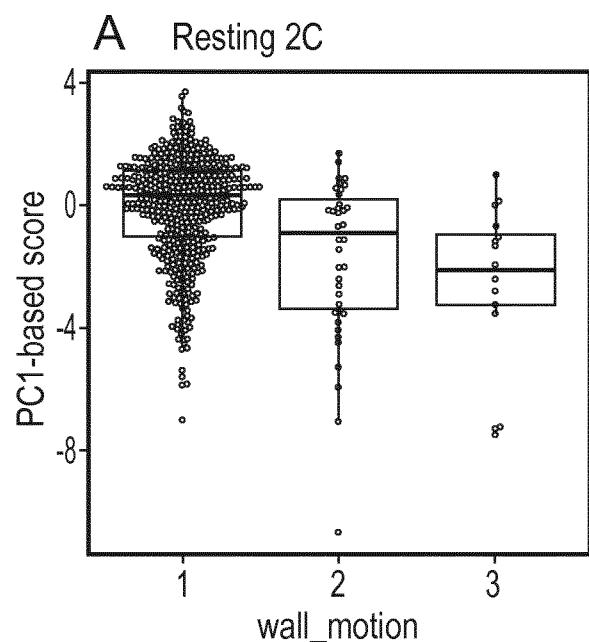


FIG. 10A

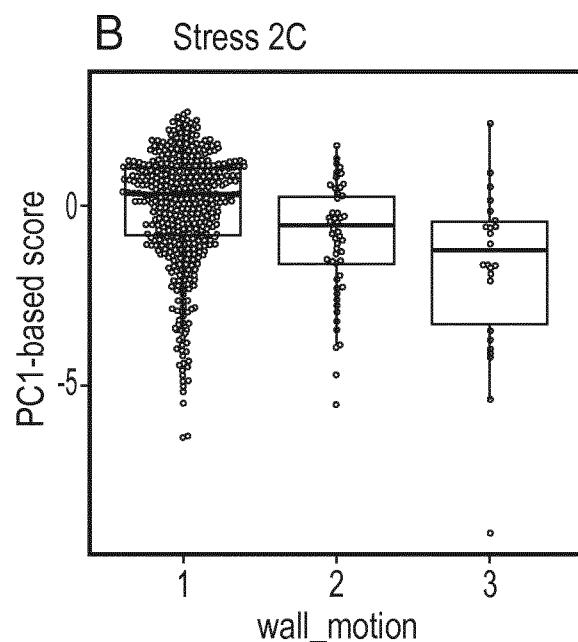


FIG. 11A

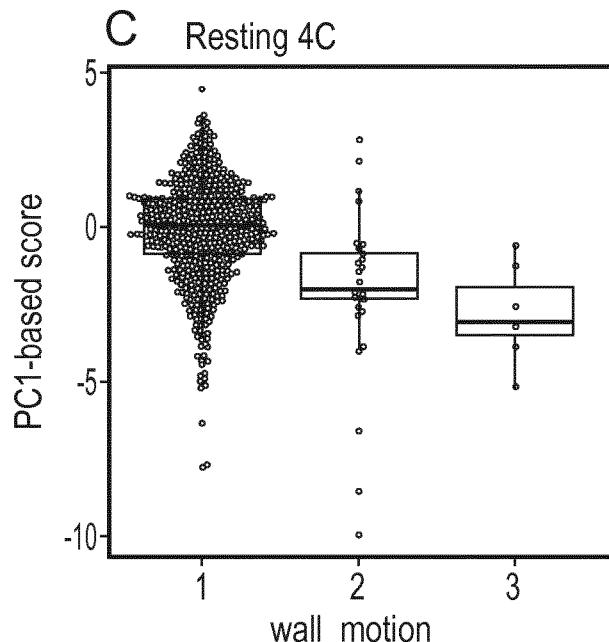


FIG. 10B

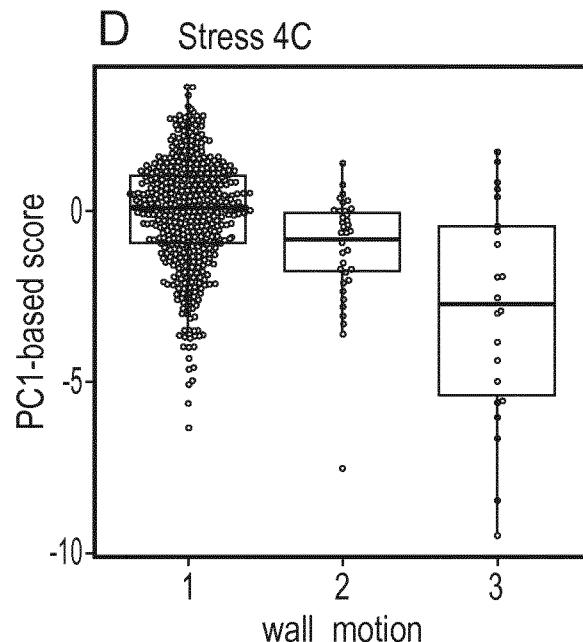


FIG. 11B

12/14

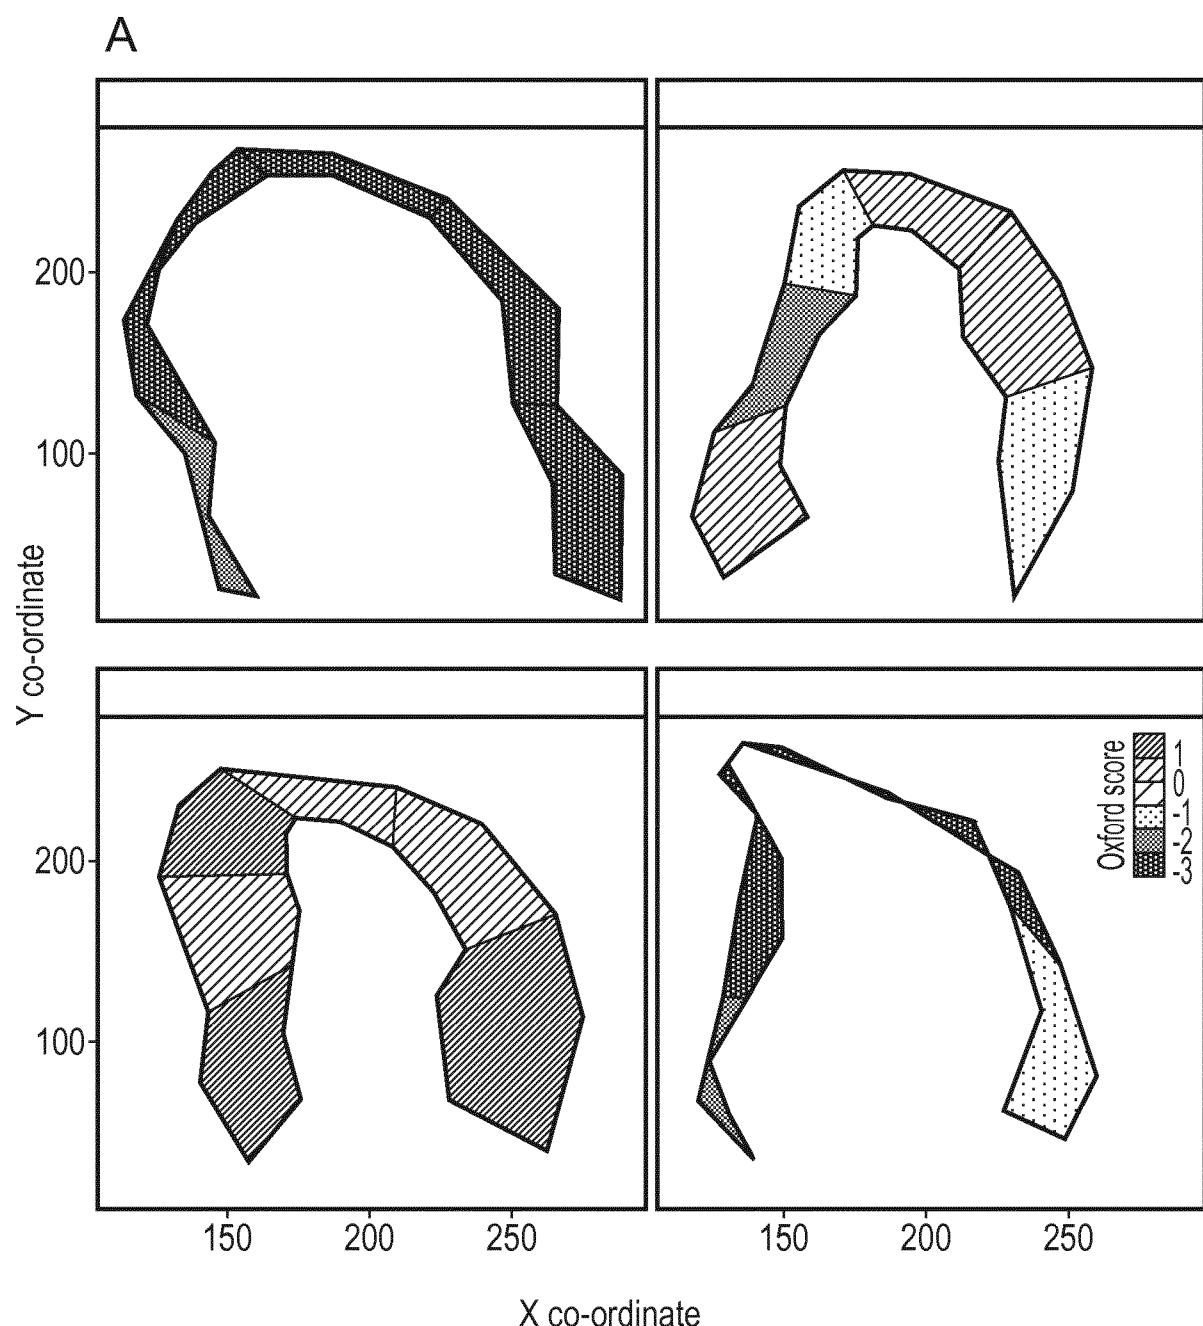
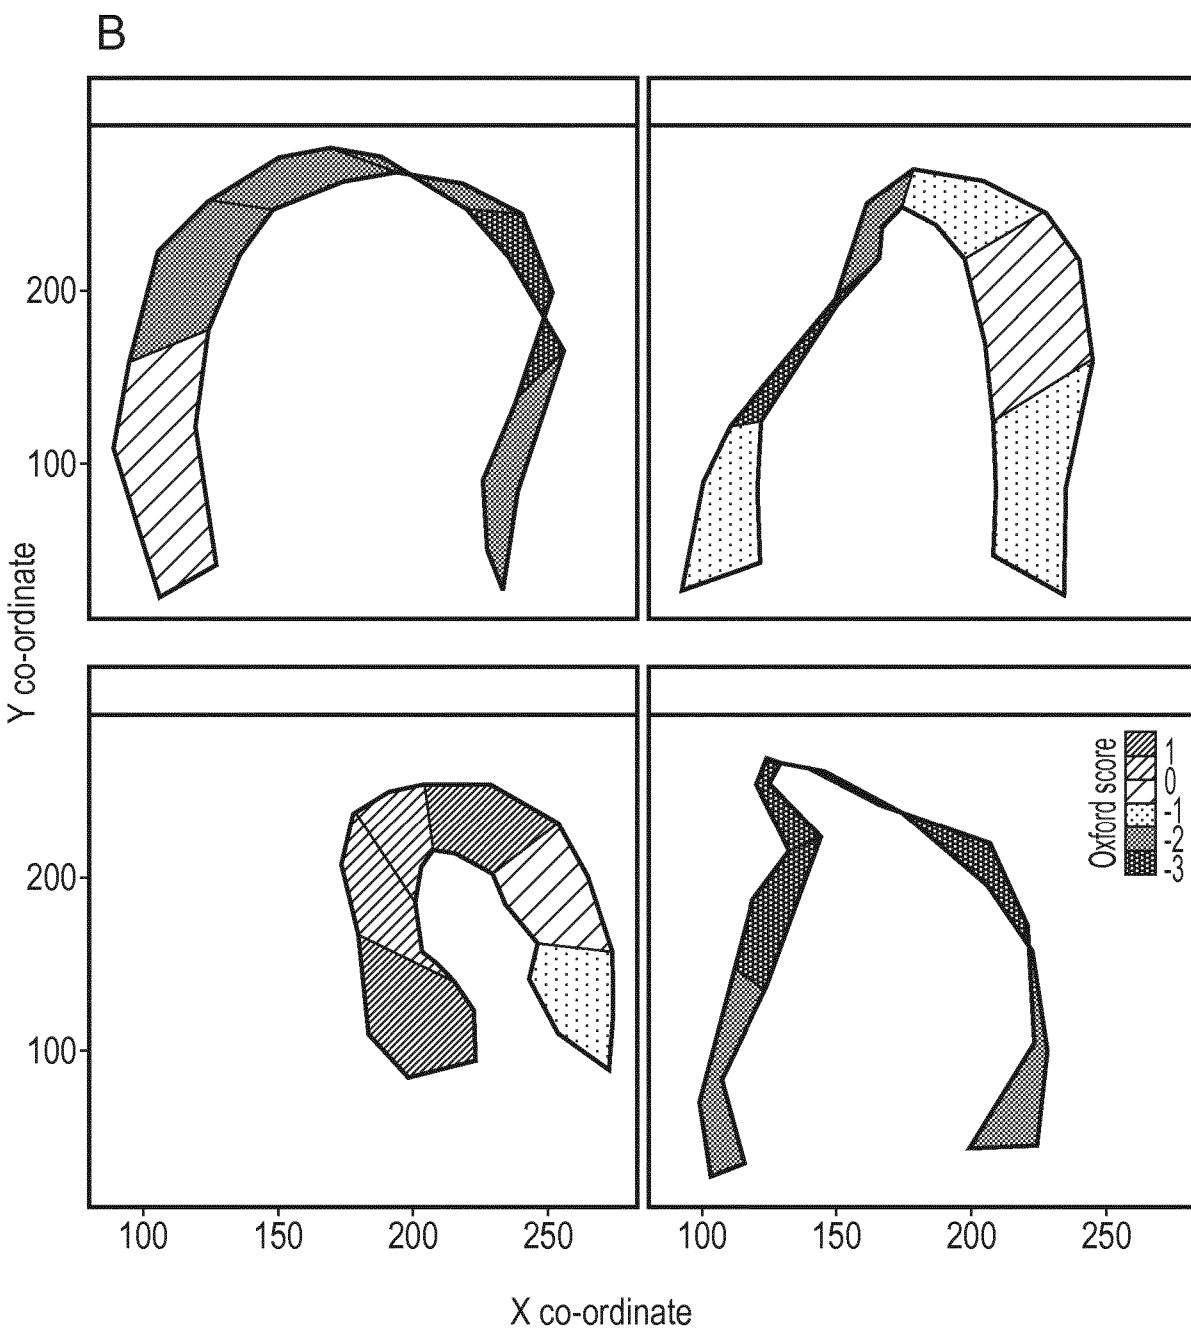


FIG. 12A

13/14



14/14

Tree View

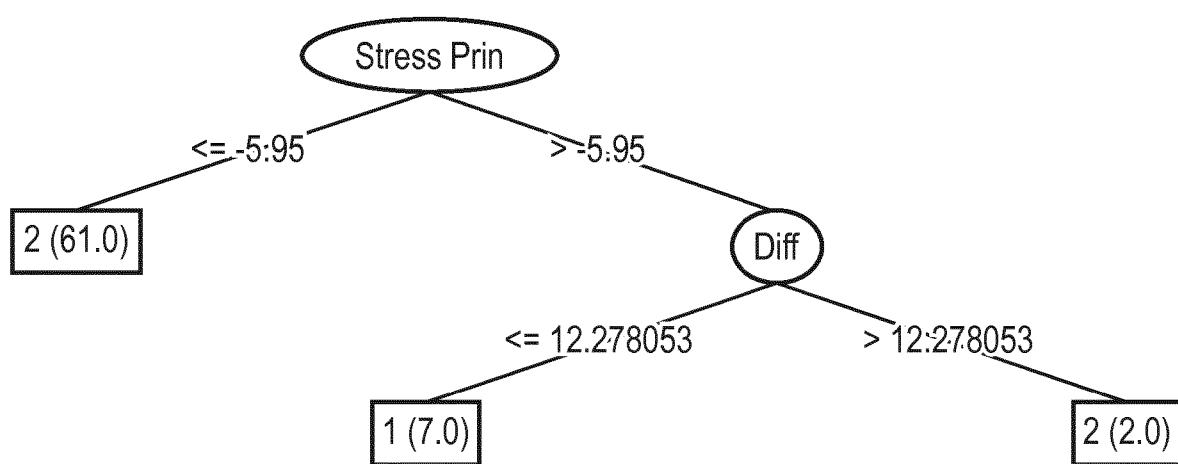


FIG. 13

A

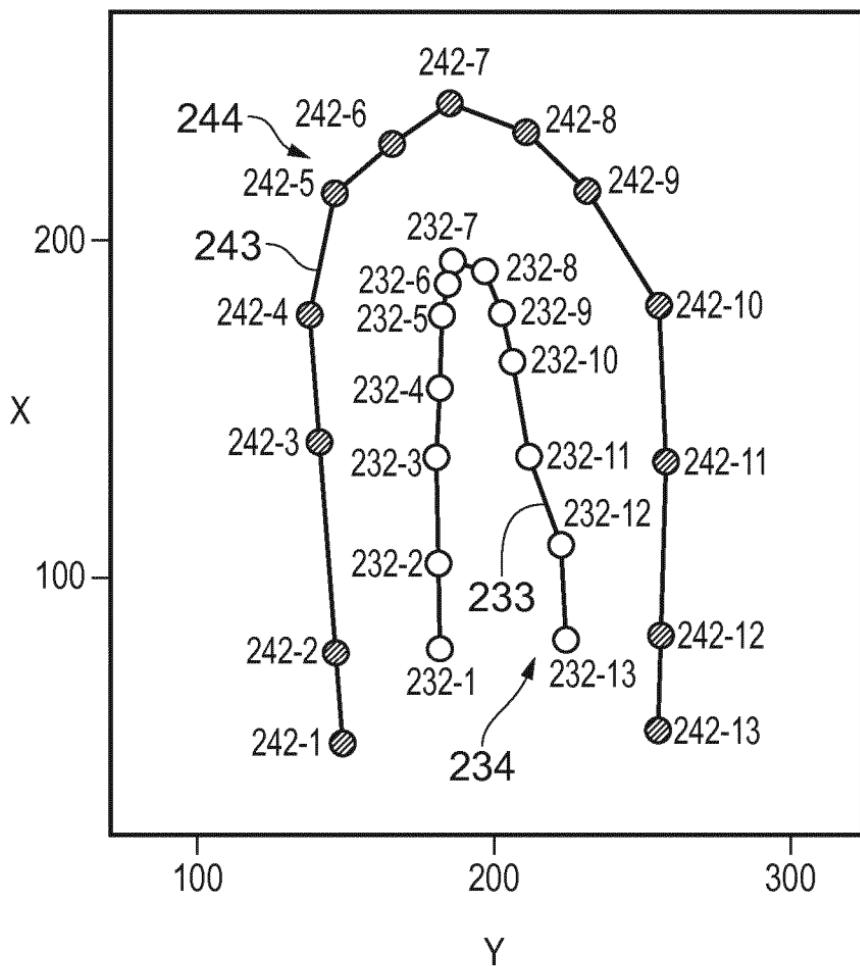


FIG. 5A