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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
24 February 2005 (24.02.2005)

PCT

(10) International Publication Number
WO 2005/017815 A2

(51) International Patent Classification[7]: G06K 9/00

(21) International Application Number:
PCT/US2004/026425

(22) International Filing Date: 13 August 2004 (13.08.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/494,646	13 August 2003 (13.08.2003)	US
10/915,076	10 August 2004 (10.08.2004)	US

(71) Applicant: **SIEMENS MEDICAL SOLUTIONS USA, INC.** [US/US]; 51 Valley Stream Parkway, Malvern, Pennsylvania 19355-1406 (US).

(72) Inventors: **KRISHNAN, Arun**; 31 Cameron Court, Exton, Pennsylvania 19341 (US). **FUNG, Glenn**; 1020 Rees Avenue, Apt. B, Bryn Mawr, Pennsylvania 19010 (US). **STOECKEL, Jonathan**; 247 Aberdeen Avenue, Exton, Pennsylvania 19341 (US).

(74) Agents: **PASCHBURG, Donald B.** et al.; Siemens Corporation- Intellectual Property Dept., 170 Wood Avenue South, Iselin, New Jersey 08830 (US).

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

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: INCORPORATING SPATIAL KNOWLEDGE FOR CLASSIFICATION

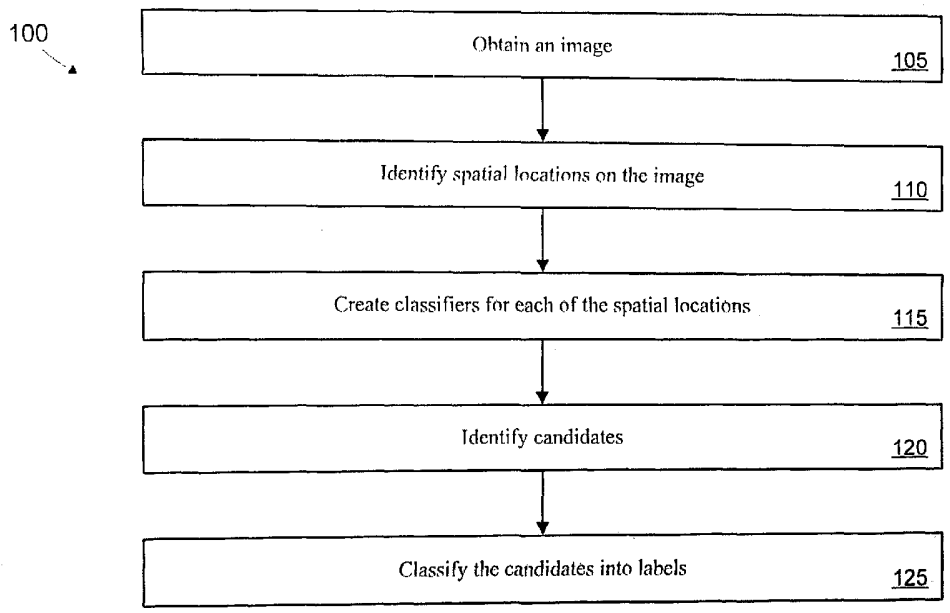

(57) Abstract: We propose using different classifiers based on the spatial location of the object. The intuitive idea behind this approach is that several classifiers may learn local concepts better than a "universal" classifier that covers the whole feature space. The use of local classifiers ensures that the objects of a particular class have a higher degree of resemblance within that particular class. The use of local classifiers also results in memory, storage and performance improvements, especially when the classifier is kernel-based. As used herein, the term "kernel-based classifier" refers to a classifier where a mapping function (i.e., the kernel) has been used to map the original training data to a higher dimensional space where the classification task may be easier.

WO 2005/017815 A2

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INCORPORATING SPATIAL KNOWLEDGE FOR CLASSIFICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority to U.S. Provisional Application No. 60/494,646, which was filed on August 13, 2003, and which is fully incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the field of machine classification, and, more particularly, to incorporating spatial knowledge for classification.

2. Description of the Related Art

A classifier is used to classify (*i.e.*, separates) objects into two or more classes. An example of a classifier is as follows. Assume we have a set, A, of objects comprising two groups (*i.e.*, classes) of the objects that we will call A+ and A–. As used herein, the term "object" refers to one or more elements in a population. The classifier, A, is a function, F, that takes every element in A and returns a label "+ "or "–", depending on what group the element is. That is, the classifier may be a FUNCTION F(A) → {–1, 1}, where –1 is a numerical value representing A– and +1 is a numerical value representing A+. The classifiers A+ and A- may represent two separate populations. For example, A+ may represent structures in the lung (*e.g.*, vessels, bronchi) and A– may represent nodules. Once the function, F, is trained from training data (*i.e.*, data with known classifications), classifications of new and unseen data can be predicted using the function, F. For example, a classifier can be trained in 10,000 known objects for which we have readings from doctors. This is commonly

referred to as a "ground truth." Based on the training from the ground truth, the classifier can be used to automatically diagnose new and unseen cases.

A conventional classifier classifies the objects into classes based on an assumption that objects of the same class have comparable feature values, that is, belong to the same distribution in the feature space. In many applications, however, objects that belong to the same class have different feature values due to, for example, their spatial location. As used herein, the term "feature" refers to one or more attributes that describe an object belonging to a particular class. For example, a nodule can be described by a vector containing a number of attributes, such as size, diameter, sphericity, etc. The vector may contain attribute values, which are termed herein as "feature values."

An existing solution to the above problem is to use a more complex classifier having a higher number of degrees of freedom. One way to create the more complex classifier is by mapping the data into a higher dimensional feature space using kernel mappings. That is, a function K (*i.e.*, the kernel) takes the original data and maps it to a higher dimensional feature space (*i.e.*, a feature space with more features) where the task of finding a classification function is easier to achieve.

For example, suppose the original data is a single case that analyzes a certain number of features. Thus, if the original data has 10 features, a vector for the original data is a 10 dimensional feature space. Suppose also that we have similar data for 1,000 other cases. We can create a similarity function that generates a value indicating how similar the original data is to each of the 1,000 other cases. Thus, a 1,000 dimensional feature space is created, mapped from the 10 dimensional feature space.

As used herein, the term "degrees of freedom" refers to the number of values in the final calculation of a statistic that are free to vary. Another way to produce the more complex classifier is to combine a number of simple classifiers, each trained using different sets of features, in series or in parallel. The results are merged to form an ensemble of classifiers.

These and other comparable approaches have the disadvantage of requiring additional training examples to compensate for the large number of degrees of freedom needed to obtain the same generalization performance. As used herein, the term "generalization performance" refers to the performance of a classifier on new and unseen data. A reduction of generalization performance may occur for many reasons. For example, in the learning (*i.e.*, training) process for classification, similar to the regression case, there is a potential risk of overfitting the training data, resulting in poor predictive performance on new and unseen cases.

SUMMARY OF THE INVENTION

In a first exemplary aspect of the present invention, a method of classifying candidates in an image is provided. The method comprises classifying a first candidate of the image using a first classifier, wherein the first classifier is trained using training data specific to the spatial location of the first candidate; and classifying a second candidate of the image using a second classifier, wherein the second classifier is trained using training data specific to the spatial location of the second candidate; wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

In a second exemplary aspect of the present invention, a method of automatically detecting a nodule on an image is provided. The method comprises obtaining an image; identifying a first spatial location and a second spatial location on

the image, wherein the first spatial location is mutually exclusive from the second spatial location; identifying a first candidate in the first spatial location and a second candidate in the second spatial location; classifying whether the first candidate is a nodule using a first classifier, wherein the first classifier is trained using training data specific to the first spatial location; and classifying whether the second candidate is a nodule using a second classifier, wherein the second classifier is trained using training data specific to the second spatial location; wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

In a third exemplary aspect of the present invention, a machine-readable medium having instructions stored thereon for execution by a processor to perform method of classifying candidates in an image is provided. The method comprises classifying a first candidate of the image using a first classifier, wherein the first classifier is trained using training data specific to the spatial location of the first candidate; and classifying a second candidate of the image using a second classifier, wherein the second classifier is trained using training data specific to the spatial location of the second candidate; wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

In a fourth exemplary embodiment of the present invention, a machine-readable medium having instructions stored thereon for execution by a processor to perform method of automatically detecting a nodule on an image is provided. The method comprises obtaining an image; identifying a first spatial location and a second spatial location on the image, wherein the first spatial location is mutually exclusive from the second spatial location; identifying a first candidate in the first spatial location and a second candidate in the second spatial location; classifying whether the first candidate is a nodule using a first classifier, wherein the first classifier is trained

using training data specific to the first spatial location; and classifying whether the second candidate is a nodule using a second classifier, wherein the second classifier is trained using training data specific to the second spatial location; wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention may be understood by reference to the following description taken in conjunction with the accompanying drawings, in which like reference numerals identify like elements, and in which:

Figure 1 depicts a method of classifying objects based on the spatial locations of the objects is illustrated, in accordance with one embodiment of the present invention; and

Figures 2A and 2B depict exemplary CT lung scans after classification of nodules, in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Illustrative embodiments of the invention are described below. In the interest of clarity, not all features of an actual implementation are described in this specification. It will of course be appreciated that in the development of any such actual embodiment, numerous implementation-specific decisions must be made to achieve the developers' specific goals, such as compliance with system-related and business-related constraints, which will vary from one implementation to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking for those of ordinary skill in the art having the benefit of this disclosure.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and are herein described in detail. It should be understood, however, that the description herein of specific embodiments is not intended to limit the invention to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

It is to be understood that the systems and methods described herein may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. In particular, at least a portion of the present invention is preferably implemented as an application comprising program instructions that are tangibly embodied on one or more program storage devices (*e.g.*, hard disk, magnetic floppy disk, RAM, ROM, CD ROM, etc.) and executable by any device or machine comprising suitable architecture, such as a general purpose digital computer having a processor, memory, and input/output interfaces. It is to be further understood that, because some of the constituent system components and process steps depicted in the accompanying Figures are preferably implemented in software, the connections between system modules (or the logic flow of method steps) may differ depending upon the manner in which the present invention is programmed. Given the teachings herein, one of ordinary skill in the related art will be able to contemplate these and similar implementations of the present invention.

Instead of using a single, complex classifier we propose using different classifiers based on the spatial location of the object. The intuitive idea behind this approach is that several classifiers may learn local concepts better than a "universal" classifier that covers the whole feature space. The use of local classifiers ensures that

the objects of a particular class have a higher degree of resemblance within that particular class. The use of local classifiers also results in memory, storage and performance improvements, especially when the classifier is kernel-based. As used herein, the term "kernel-based classifier" refers to a classifier where a mapping function (*i.e.*, the kernel) has been used to map the original training data to a higher dimensional space where the classification task may be easier.

Referring now to Figure 1, a method 100 of classifying objects based on the spatial locations of the objects is illustrated, in accordance with one embodiment of the present invention. The method 100 obtains (at 105) an image. In one embodiment, the image may be obtained from a physical scan of the human body, such as from a computed tomography ("CT"), magnetic resonance imaging ("MRI") or ultrasound device. In another embodiment, the image may be obtained from a satellite imaging system. It should be appreciated that any of a variety of methods and apparatus may be used to obtain the image, as contemplated by those skilled in the art.

The method 100 identifies (at 110) spatial locations for elements on the image. In an image of a lung, for example, potential locations may include a wall-attached location representing elements attached to the lung wall, a bronchi classifier representing elements attached to the bronchi, and an inside-lung classifier representing elements inside the lung. Other images of the human body can be used, such as the heart. For another example, a satellite image may be classified into a type of landmass (*e.g.*, mountain, desert, rainforest). The spatial locations may be identified through, for example, a graphical user-interface. The spatial locations may also be identified automatically using, for example, previous knowledge of various locations on an image. It should be appreciated that various other automated methods

(*e.g.*, pattern matching) for identifying spatial locations may be used, as contemplated by those skilled in the art.

Classifiers are created (at 115) for each of the spatial locations. Using the lung scan example above, potential classifiers may include a wall-attached classifier representing elements attached to the lung wall, a bronchi classifier representing elements attached to the bronchi, and an inside-lung classifier representing elements inside the lung. The classifiers may be trained using previously-known data. For example, using the lung scan example above, the classifiers may be trained using training data of elements specified for the particular spatial location. Training data may comprise known and previously-verified data. For example, doctor-verified cases determining whether a nodule exists can be used as training data. The various cases may be separated into spatial locations corresponding to the different classifiers for training purposes. Thus, a wall-attached classifier would be trained using only training data of elements attached to the lung wall. The advantage of such a method is that the classifier takes advantage of any novelties (known and unknown) of elements in a particular location. For example, classifying elements in the bronchi may be different from classifying elements inside the lung.

The method 100 identifies (at 120) candidates. As used herein, the term "candidates" refers to elements (*i.e.*, structures) of interest in the image. Because the classifiers are unique to the spatial locations, the classification of candidates will be more precise than not taking into account the spatial location. Further, using such "local classifiers" also result in memory, storage and performance improvements over complex, universal classifiers. Candidates may be identified manually or automatically using any of a variety of methods and apparatus known to those skilled

in the art. For example, computer-aided automatic identification of suspicious regions of interest is well-known in the field of cancer screening.

The method 100 classifies (at 125) the candidates into labels using the appropriate classifier associated with the spatial location where the candidate is located. For example, in the lung scan example above, if a candidate is located inside the lung, then the candidate is classified using the inside-lung classifier. The candidate may be labeled "nodule" representing that a structure (*i.e.*, the element) in the image is a nodule and "no-nodule" representing that the structure is not a nodule. It should be appreciated that any of a variety of labels may be used, as contemplated by those skilled in the art. Further, a candidate may be classified with more than two labels.

It should be appreciated that particular labels can be physically marked on the image for easy identification. For example, in the lung scan example above, a candidate identified as a nodule may be marked so that a human observer (*e.g.*, a doctor) can easily identify the pathology. Different marking types and shapes may be used to differentiate between certain classifiers, as is contemplated by those skilled in the art. An example of a marking may be a square or a circle surrounding each identified nodule.

Figures 2A and 2B show exemplary CT lung scans that are marked after classification, in accordance with one embodiment of the present invention. The square markings refer to wall-attached nodules, and the circle markings refer to inside-lung nodules.

The particular embodiments disclosed above are illustrative only, as the invention may be modified and practiced in different but equivalent manners apparent to those skilled in the art having the benefit of the teachings herein. Furthermore, no

limitations are intended to the details of construction or design herein shown, other than as described in the claims below. It is therefore evident that the particular embodiments disclosed above may be altered or modified and all such variations are considered within the scope and spirit of the invention. Accordingly, the protection sought herein is as set forth in the claims below.

WHAT IS CLAIMED IS:

1. A method of classifying candidates in an image, comprising:

classifying a first candidate of the image using a first classifier, wherein the first classifier is trained using training data specific to the spatial location of the first candidate; and

classifying a second candidate of the image using a second classifier, wherein the second classifier is trained using training data specific to the spatial location of the second candidate;

wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

2. The method of claim 1, further comprising obtaining the image.

3. The method of claim 2, wherein obtaining the image comprises obtaining at least one of a CT, MRI or ultrasound image of at least part of a human body.

4. The method of claim 2, wherein obtaining the image comprises obtaining a satellite image.

5. The method of claim 1, further comprising identifying the spatial locations on the image.

6. The method of claim 5, wherein identifying the spatial locations on the image comprises providing a graphical user-interface for manually identifying the spatial locations on the image.

7. The method of claim 5, wherein identifying the spatial locations on the image comprises automatically identifying the spatial locations on the image.

8. The method of claim 7, wherein automatically identifying the spatial locations on the image comprises automatically identifying the spatial locations on the image using known information of the spatial locations.

9. The method of claim 7, wherein automatically identifying the spatial locations on the image comprises automatically identifying the spatial locations on the image using pattern matching.

10. The method of claim 5, wherein identifying the spatial locations on the image of at least part of a human body.

11. The method of claim 10, wherein identifying the spatial locations on the image of at least part of a human body comprises identifying the spatial locations on an image of a lung.

12. The method of claim 11, wherein the spatial locations on an image of a lung comprise a wall-attached area, an inside-lung area, and a bronchi area.

13. The method of claim 10, wherein identifying the spatial locations on the image of at least part of a human body comprises identifying the spatial locations on an image of a heart.

14. The method of claim 1, further comprising identifying the first candidate and the second candidate on the image.

15. The method of claim 14, wherein identifying the first candidate and the second candidate on the image comprises providing a graphical user-interface for manually identifying the first candidate and the second candidate on the image.

16. The method of claim 14, wherein identifying the first candidate and the second candidate on the image comprises automatically identifying the first candidate and the second candidate on the image.

17. The method of claim 16, wherein automatically identifying the first candidate and the second candidate on the image comprises automatically identifying the first candidate and the second candidate on the image using pattern matching.

18. The method of claim 1, wherein classifying a first candidate comprises automatically classifying whether the first candidate is a nodule, and wherein classifying a second candidate comprises automatically classifying whether a second candidate is a nodule.

19. The method of claim 1, wherein the training data comprises doctor-verified data.

20. A method of automatically detecting a nodule on an image, comprising:

obtaining an image;

identifying a first spatial location and a second spatial location on the image, wherein the first spatial location is mutually exclusive from the second spatial location;

identifying a first candidate in the first spatial location and a second candidate in the second spatial location;

classifying whether the first candidate is a nodule using a first classifier, wherein the first classifier is trained using training data specific to the first spatial location; and

classifying whether the second candidate is a nodule using a second classifier, wherein the second classifier is trained using training data specific to the second spatial location;

wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

21. The method of claim 20, further comprising:

marking the first candidate if the first candidate is a nodule; and

marking the second candidate if the second candidate is a nodule.

22. A machine-readable medium having instructions stored thereon for execution by a processor to perform method of classifying candidates in an image, the method comprising:

classifying a first candidate of the image using a first classifier, wherein the first classifier is trained using training data specific to the spatial location of the first candidate; and

classifying a second candidate of the image using a second classifier, wherein the second classifier is trained using training data specific to the spatial location of the second candidate;

wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

23. A machine-readable medium having instructions stored thereon for execution by a processor to perform method of automatically detecting a nodule on an image, the method comprising:

obtaining an image;

identifying a first spatial location and a second spatial location on the image, wherein the first spatial location is mutually exclusive from the second spatial location;

identifying a first candidate in the first spatial location and a second candidate in the second spatial location;

classifying whether the first candidate is a nodule using a first classifier, wherein the first classifier is trained using training data specific to the first spatial location; and

classifying whether the second candidate is a nodule using a second classifier, wherein the second classifier is trained using training data specific to the second spatial location;

wherein the spatial location of the first candidate is mutually exclusive from the spatial location of the second candidate.

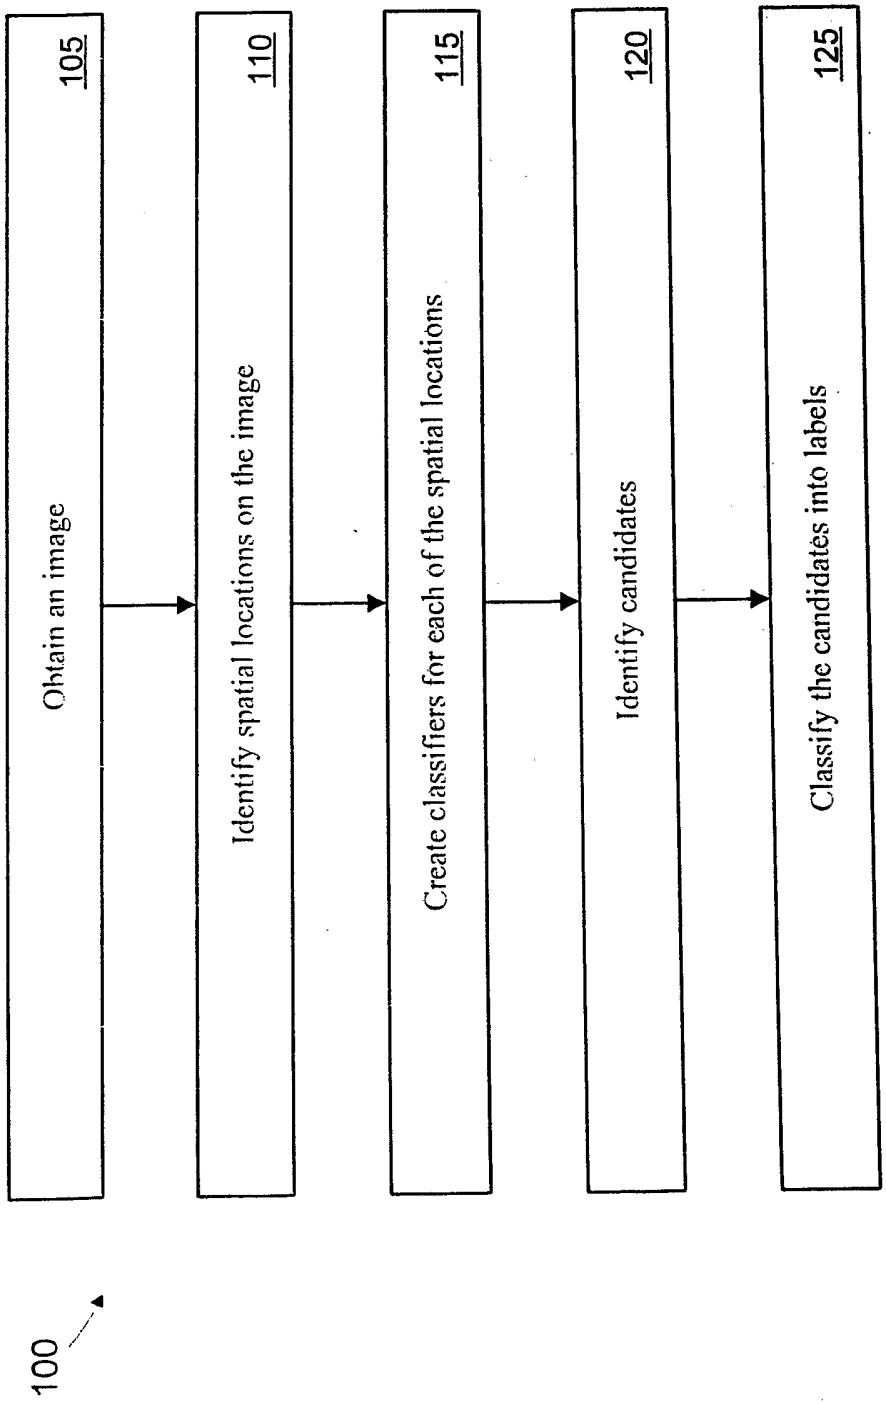
Obtain an image
105
Identify spatial locations on the image
110
Create classifiers for each of the spatial locations
115
Identify candidates
120
Classify the candidates into labels
125
100

Figure 1

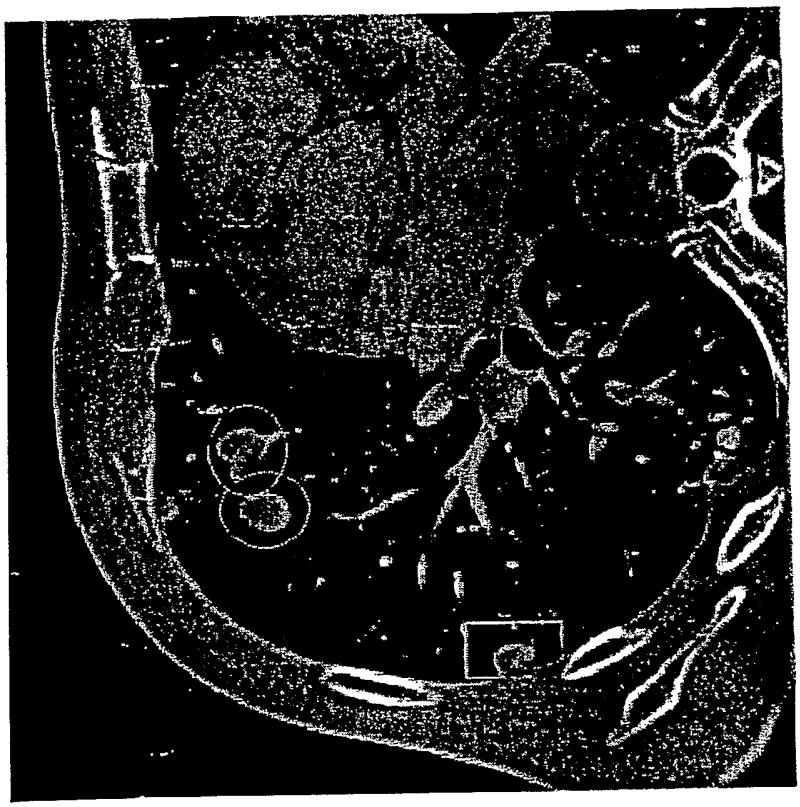

Figure 2B

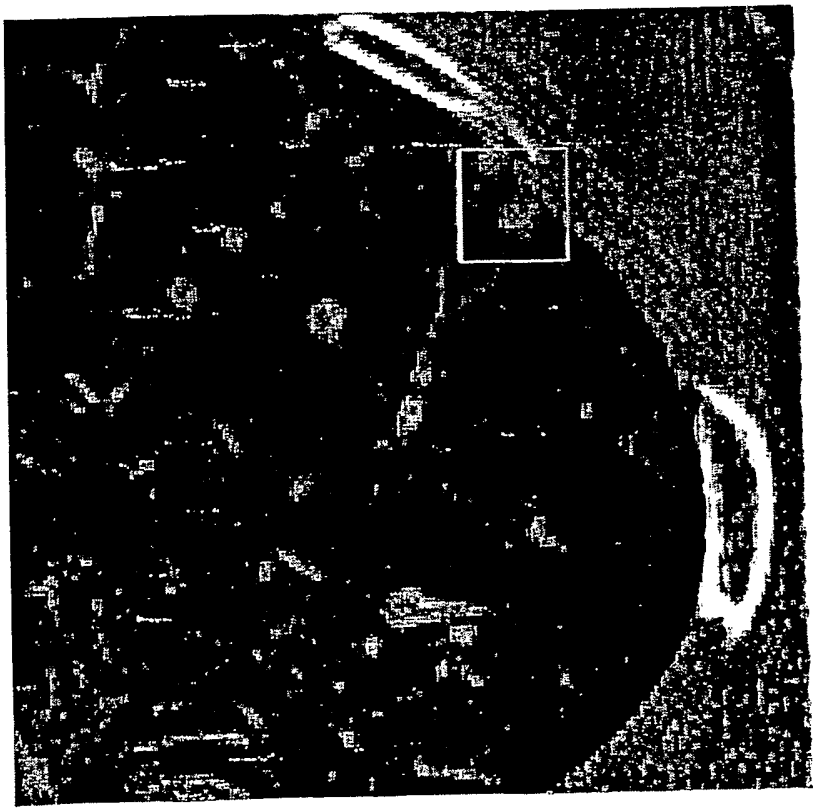

Figure 2A