

(19) World Intellectual Property Organization
International Bureau



PCT



(43) International Publication Date
31 December 2008 (31.12.2008)

(10) International Publication Number
WO 2009/003198 A1

(51) International Patent Classification:
G06K 9/00 (2006.01)

(21) International Application Number:
PCT/US2008/068826

(22) International Filing Date: 30 June 2008 (30.06.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/946,814 28 June 2007 (28.06.2007) US

(71) Applicant (for all designated States except US): **MAYO FOUNDATION FOR MEDICAL EDUCATION AND RESEARCH** [US/US]; 200 First Street SE, Rochester, MN 55905 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **PATRIARCHE, Julia, W.** [CA/US]; 300 16th Street NE #302, Rochester, MN 55906 (US). **ERICKSON, Bradley, J.** [US/US]; 730 Fox Chase Road SW, Rochester, MN 55902 (US).

(74) Agent: **COOK, Jack, M.**; Quarles & Brady LLP, 411 E. Wisconsin Ave., Milwaukee, WI 53202 (US).

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

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

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

(54) Title: SYSTEM AND METHOD FOR AUTOMATICALLY DETECTING CHANGE IN A SERIES OF MEDICAL IMAGES OF A SUBJECT OVER TIME




FIG. 18

(57) Abstract: A system and method for detecting changes in a series of images medical includes an automated sample point generator that reviews the series of medical images and automatically detect portions of the series of medical images indicative of a material expected to be in a region of interest (ROI) and designate sample points therefrom. A normalizer is configured to utilize the sample points to normalize intensities within the series of medical images and an image registration system is configured to utilize the sample points to automatically register the series of images such that common areas within the ROI of the subject are aligned. A change tracker provided to automatically identify changes between the series of medical images and display an indication of the identified changes.

WO 2009/003198 A1

SYSTEM AND METHOD FOR AUTOMATICALLY DETECTING CHANGE IN A SERIES OF MEDICAL IMAGES OF A SUBJECT OVER TIME

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is based on U.S. Provisional Patent Application Serial No. 60/946,814 filed on June 28, 2007, and entitled "AUTOMATED DETECTION OF CHANGE IN SERIAL IMAGING STUDIES OF THE BRAIN."

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

[0002] This invention was made with government support under Grant No. T32 NS07494 and T15 LMO7041-23 awarded by the National Institute of Health. The United States Government has certain rights in this invention.

BACKGROUND OF THE INVENTION

[0003] The field of the invention is nuclear magnetic resonance imaging methods and systems. More particularly, the invention relates to a system and method for tracking physical changes in a series of medical images of a given patient over time.

[0004] When a substance such as human tissue is subjected to a uniform magnetic field (polarizing field B_0), the individual magnetic moments of the spins in the tissue attempt to align with this polarizing field, but precess about it in random order at their characteristic Larmor frequency. If the substance, or tissue, is subjected to a magnetic field (excitation field B_1) that is in the x-y plane and that is near the Larmor frequency, the net aligned moment, M_z , may be rotated, or "tipped", into the x-y plane to produce a net transverse magnetic moment M_t . A signal is emitted by the excited spins after the excitation signal B_1 is terminated, this signal may be received and processed to form an image.

[0005] When utilizing these signals to produce images, magnetic field gradients (G_x , G_y and G_z) are employed. Typically, the region to be imaged is scanned by a sequence of measurement cycles in which these gradients vary according to the particular localization method being used. The resulting set of received NMR signals are digitized and processed to reconstruct the image using one of many well known reconstruction techniques.

[0006] An MRI system may be used to acquire many types of images from a particular anatomical structure, for example, the brain. Such images employ contrast mechanisms that enable different brain tissues and lesions to be identified. The usual practice is to acquire a set of MRI images and then manually segment different

tissues and lesions and provide sample points of the different tissues for use in subsequent processing. Only limited attempts have been made to implement automated sample point generation algorithms. Some classification algorithms provide very rudimentary automated sample points generation, with little use of domain specific knowledge (e.g., fuzzy k-means), where the user supplies the number of pure tissues and the algorithm attempts to locate the cluster centroids by examining only the image intensities.

[0007] Atlas based segmentation methods can be used as one method for the automatic generation of samples, and there have been examples of such implementations in the literature (e.g., Linguraru et al., 2006). However, by definition these atlas based algorithms make assumptions regarding anatomy, which may not be true considering pathology and surgical interventions. One important motivation for the performance of anatomical magnetic resonance imaging and image analysis is the examination of pathology (e.g. tumors, multiple sclerosis, etc.). In such cases, pathology may be quite wide-spread, and may present in a large variety of ways, causing deviations from "normal" anatomy. Furthermore, resections may be performed in such cases, as well, causing the particular patient's brain to deviate even further from the atlas-based norms.

[0008] Classification of lesions, per se, is not new, and many solutions to classify multi-spectral data exist, both within and outside of the domain of medical imaging. Some classification methods are supervised and require that the user supply sample data. Some examples of supervised classification algorithms are thresholding, Euclidean distance, Mahalanobis distance, K nearest neighbor, Bayesian, and neural network. Other classification algorithms are unsupervised and require no sample data. Some examples of unsupervised classification algorithms include chain method, and k-means. It should be noted, however, that unsupervised classification algorithms typically require some rudimentary data be supplied by the user. For example, the chain method requires that a threshold be supplied, and k-means requires that the actual number of clusters be supplied. Some classification algorithms categorize each data point, that is, each voxel in the case of image data, into one of a number of discrete categories. Other classification algorithms allow partial membership in multiple categories. Still, other classification algorithms allow partial membership in multiple classes initially, but then defuzzify the membership data into discrete categories in a later step.

[0009] A large fraction of existing classification algorithms are devoted to correctly assigning voxels into discrete categories. In fact, many authors write explicitly that the purpose of classification is to assign multispectral data points into discrete categories and some then add as an after-thought that the assignment may also be made fuzzy. Furthermore, even the algorithms that purport to possess fuzzy classification capability, still have a stated purpose of achieving the highest accuracy with voxels that do in fact belong purely to one category or another, such as. Mahalanobis distance based classifiers. These algorithms acknowledge that in their intended domain, the overwhelming majority of voxels will, in fact, belong purely to one category or another and; therefore, in order to achieve the highest average accuracy, these methods use membership functions that achieve the highest accuracy for voxels that are in fact purely one category. However, this is achieved at the expense of accuracy for voxels that do in fact possess mixed membership.

[0010] In the case of lesions of the white matter of the brain, the voxels that, in fact, possess partial membership are the voxels that are of greatest interest. There are a variety of reasons for an algorithm to be developed that is focused on accurately quantifying partial membership and partial character, rather than categorizing voxels into discrete categories. For example, in the case of pathology (lesion and enhancing lesion), voxels are rarely fully abnormal. Rather, in such cases, abnormality exists in degree. It is subtly abnormal voxels for which computational assistance might find a high degree of usefulness, because these subtly abnormal voxels are relatively difficult to see with the unaided eye. Furthermore, accurate quantification of these partial volume and partial membership characteristics is important, because it has been shown that, at least in tumors, the rate of growth is suggestive of prognosis. Thus, particularly for small lesions, the ability to accurately quantify these effects places a distinct theoretical lower bound on the size of the lesions from which prognosis may be quantified.

[0011] One group of authors created a classification procedure specifically geared for recognizing lesions that recognized the importance of preserving partial volume effects via a feature extraction step involving a linear combination of the original images (Hamid Soltanian-Zadeh, 1998). These authors did not, however, design their algorithm with the real behavior of lesion intensities in feature space in mind. Rather, the algorithm, which used an Eigenimage feature extraction step, was more tuned to the behavior of noise and contrast, than the lesions themselves. The algorithm required that the images demonstrate frank lesion and that manual sample

points be supplied, which required an investment of time by the user and that introduced variability into the process. Furthermore, the authors used thresholding as their noise reduction scheme.

[0012] It has been recognized for a considerable period of time that it is challenging for a human observer to compare images side by side that are nearly the same, in order to identify the differences. This problem exists in a large number of fields of imaging, both within and outside of, medicine. Some examples of non-medical change detection topics include remote sensing, astronomy, surveillance, geology, reconnaissance photography, and the like. The development of computer algorithms to compare serial images has been a fruitful field of study in a variety of areas. Probably the most work has been done in the field of remote sensing, where satellite imagery of the earth, using various kinds of sensors, has been used in an attempt to detect various kinds of changes. In this sub-discipline of change detection as in others, imagery is to be compared between image sets acquired at two time-points. The various sub-fields of change detection have important points in common. First, the motivation for change detection is generally the same in all cases. Side-by-side presentation of images for the purposes of identifying small differences is a sub-optimal method of display, because it requires the viewer to sequentially examine each area and each feature to be compared. Second, imaging systems, which include sensors, data storage facilities, and the like, are becoming more and more prevalent and more and more capable of storing vast amounts of data. This phenomena has been called "information overload" and there is little doubt that this trend will continue. With vast amounts of data being collected, it will be unreasonable or impossible for a human to make every imaginable comparison, but it is not at all unreasonable for this task to be performed by computer.

[0013] There are a variety of reasons that the comparison of serial images is challenging. One reason that the problem of change detection is difficult is that side-by-side presentation of images is poorly matched to our visual systems. That is, the particular workings of our visual systems means that observers much sequentially examine each feature of interest, and explicitly compare them. This is time consuming, particularly as the data sets and features of interest become large, and it additionally biases us against detecting changes that we do not expect, but which nonetheless may still be relevant and important. Another reason that change detection is challenging is that the disease related changes of interest are likely to be confounded with acquisition related changes, which are not of interest. Still another

reason is that comparisons that might be desired are more complicated than simple comparison of the intensities in the images. With greater and greater complexity of the questions to be asked and the comparisons to be made, the task becomes more and more difficult for an unaided observer. Yet another reason the task of comparison of images is difficult is, as mentioned above, information overload. Information overload is a problem that will almost definitely increase as the sophistication of imaging hardware improves.

[0014] In its simplest form, change detection may be performed as subtraction of two serial images as shown in Fig. 1. This approach has been used with great effectiveness. More sophisticated approaches may be imagined and have been used, where some processing is performed on the images, to produce derivative data that is then compared as shown in Fig. 2. Spatial registration is a simple example of such processing, which is commonly applied in remote sensing, and medical image analysis. A large variety of image processing steps exist, and can be conceived of. These may require multiple steps in a sequence, for example, inhomogeneity correction followed by registration, followed some form of image understanding, for example, classification in remote sensing or medical imaging, followed by comparison of the classified images. In the case of surveillance, where it could, for example, be desired to observe when a new person has entered the visual field of the camera, face recognition may be one of the steps. Typically, the more sophisticated processing steps apply domain specific knowledge. The domain specific knowledge, which is used to improve the performance of change detection algorithms and to attune the algorithm to the detection of specific kinds of changes of interest, may be very sophisticated and varies greatly from application to application.

[0015] However, the behavior of the sensors, the meaning of intensities, the subject being imaged, and the like vary greatly from domain to domain and; thus, while papers in the field of remote sensing may compare log-ratio images, these may not make any sense at all in the context of comparison of serial MR brain studies. Likewise, even within brain MRI, the subject of change detection is made broad by the different kinds of domain-specific knowledge that may be applied, and the kind of derivative information that may be generated and compared from one acquisition to another. One set of analyses might be used to detect and characterize changes in white matter lesions, while a completely different set of analyses may be used to detect and characterize changes due to atrophy, for example, identification of discrete boundaries using finite element models with sub-voxel resolution, and the

comparison of the position of these boundaries from one acquisition to another. In fact, there exist a large range of possible formulations of these post-acquisition processing steps and inter time-point comparisons.

[0016] Within the context of change detection in brain MRI of white matter pathology, there are a variety of causes of inter-observer variability. One is that there is presently no objective definition of stability or progression. Additionally, changes that are relatively subtle, either in extent or degree, may be difficult to see. This is especially true when the images are displayed side-by-side, not registered, with intensity inhomogeneities and the like. Manually identifying changes in such a context requires the observer to mentally deconfound the data, which is challenging for a human observer. It can be imagined that a computer might be very well suited to helping with this task, since, theoretically, computers have unlimited memory, which is in stark contrast with human visual memory and short-term memory, both of which are limited. Thus, computers can apply a theoretically unlimited number of intermediate processing steps. Changes might be spread across multiple pulse sequences, and they might be indefinite in one slice, but more definite in consideration of multiple slices. This, however, requires assimilation and integration of a number of slices, before reaching a decision. Such a process is not natively simple for the human visual system, at least when the images are displayed as slices. However, such a process is much simpler for a computer.

[0017] There are other “expectation-oriented” reasons that change detection is challenging. Changes that occur in unexpected locations or that are unexpected in terms of their character, may be missed. There are a variety of other issues that might cause a neuroradiologist to miss changes, such as “satisfaction of search,” where a radiologist stops looking when they find imaging features that explain the symptoms motivating the scan in the first place; however, there may still be other important findings in the images. Sometimes, “information overload” makes change detection difficult, due simply to the sheer volume of data presented. This problem will almost definitely increase as scanners produce greater and greater quantities of data. In a general sense, computers excel at methodically wading through large amounts of data, looking for needles in hay-stacks, integrating large amounts of data at once, applying serial processing steps, analyzing data mathematically, and bringing noteworthy observations to the attention of the neuroradiologist.

[0018] The detection of regions of signal embedded in zero-mean background noise is a recurring problem in various fields of medical imaging, including

classification, fMRI, and change detection. The detection of such signals is also an important problem outside of medical imaging in such fields as remote sensing, surveillance, astronomy, and geology, to name a few. A variety of approaches have been taken to identify regions of signal indicating a desired objects embedded in regions of noise in imagery data. By far the most common method that has been used is simple thresholding, where pixels or voxels possessing image intensities lower than some threshold value are set to 0, while image intensities larger than the threshold value are retained. That is, image intensities below the threshold are identified as pixels or voxels are considered to not belong to the object and those above the threshold are considered to belong to the object. Thresholding is limited in that it does not allow simultaneous rejection of noise concurrent with retention of low intensity object voxels. One attempt to approach this limitation is commonly used in fMRI circles. It involves the use of a "smearing" filter on the image data prior to thresholding. Essentially, this approach has two effects. First, it smears the voxel intensities of regions containing actual objects into their neighbors. If the regions of actual objects contain some high intensity voxels, these high intensities can reinforce their low intensity neighbors, which then become included in the region retained by thresholding. The second effect of these smearing filters is for regions not containing actual objects, which presumably contain zero-mean noise, to diminish in intensity. In many cases, this increases the degree to which noisy regions are rejected by the process of thresholding. Although popular in some circles, this approach is in many ways intuitively unsatisfactory. This is primarily because it works most effectively when the actual object possesses high intensity, which is certainly not always the case, rather, it is very frequently desired to detect subtle objects. This is also because the method virtually guarantees the erroneous detection of a penumbra of noise voxels around an intense object. Furthermore, such a method can result in false negatives. That is, if dark voxels around a moderately bright object are smeared into the moderately bright object, the intensities of the voxels making up the moderately bright object to drop below the threshold.

[0019] Another filtering approach that has been used involves the application of anisotropic diffusion filters to effect the blurring. Anisotropic diffusion filters are somewhat like low pass filters, except that the blurring is reduced in regions where the gradient of the voxel intensities is high, for example, edges. In cases where objects possess edges, for example, discrete anatomical structures, such an

approach might be helpful. However, there are cases where the objects of interest do not possess definite edges at their periphery or where the objects themselves are heterogeneous internally. In the case of objects with very low mean intensities, with mean intensities close to zero and intensity distributions close to or within the level of noise, any gradients surrounding the objects will almost surely exist at a very low level. In each of these cases, an anisotropic diffusion filter would demonstrate limited effectiveness. On the positive side, an anisotropic diffusion filter might be expected to help to reduce the zero mean noise of the background, and in that respect it might aid in the use of simple thresholding to accomplish noise rejection. However, this approach significantly reduces the background noise using an anisotropic diffusion filter that may require the application of multiple iterations of the filter, which causes a significant penumbra of falsely detected voxels around the actual objects. Furthermore, anisotropic diffusion filters are highly parameter dependent, which is inconsistent with applications that desire automation and robustness. Some groups have used edge-finding algorithms themselves with no blurring to identify objects; however, as with anisotropic diffusion filters, these methods do not work when the objects do not possess external edges.

[0020] Fuzzy connectedness is a method that has been used to identify objects in images. In practical terms, fuzzy connectedness requires sample points, which is also inconsistent with systems that desire automation. Fuzzy connectedness still hinges on the memberships of single voxels, which is weakest link in the chain, in the terms of the creators of the method. The goal is to identify regions of voxels potentially possessing intensities within the level of the noise floor. By emphasizing the use of chains of connectedness, fuzzy connectedness does not simultaneously emphasize the relative positions of the neighbors with respect to trial voxels and relative to one another. This means that fuzzy connectedness does not make use of the *a priori* knowledge that one trait of real-world objects is spatial continuity. For example, real world regions or objects in brain MR images tend to be fat, at least compared to a long, thin, winding chain, and not full of holes. Fuzzy connectedness also does not provide the ability to weigh the size of a region against the memberships of its constituent voxels.

[0021] Another prior method in which voxels are accepted as belonging to objects and not belonging to background noise, determines not only whether they exceed a threshold, but some number of their neighbors also exceed the threshold. This method is interesting in as much as it is an initial attempt to balance spatial extent

against magnitude of voxel value, in order to establish whether or not a voxel belongs to a real underlying object or whether its intensity is due only to noise. Another method uses fixed dimension kernels, for example, 3x3x3 voxels, and applies a test of significance to help reduce false positives. This method is interesting for the same reasons as the prior method, but this method additionally enforces a higher degree of spatial continuousness. The restricted shape and extent of the regions used by this method, however, make it un-ideal. Another method, demonstrated in the context of surveillance, divides the overall image area into fixed test areas, for example, 4x3 voxels in extent, and applies formal statistical tests to compare the intensity distributions contained by the regions at two time-points, in order to establish whether the contents of each region are the same or different from one time point to the next. This is an interesting approach, for its use of formal statistical tests to decide whether a change had occurred or not. As above, however, the use of fixed position and size test areas is unnecessarily limiting.

[0022] Therefore, it would be desirable to have a system and method for accurately and consistently analyzing medical images that does not fall prey to the above-discussed drawbacks.

SUMMARY OF THE INVENTION

[0023] The present invention overcomes the aforementioned drawbacks by providing a system for detecting changes in medical images, for example, changes in brain lesions as measured by a series of MRI images. More particularly, it includes a sample point generator that identifies sample pixels in MRI images that represent different types of brain tissues. The system also includes a lesion detector that identifies and quantifies abnormal tissues using the sample pixels. A lesion change detector compares lesions identified currently with previously identified lesions. In addition, an object boundary identification process may be used with the system to reduce the effects of noise and more definitively identify lesion boundaries and changes.

[0024] In accordance with one aspect of the present invention, a method is disclosed that automatically, repeatably, and reliably generates sample points from brain MRI images, which are suitable for use as input to a variety of image processing methods. The method is designed to use simple and reliable knowledge regarding the brain and magnetic resonance images, so that it functions normally regardless of pathology, interventions, or particulars of the acquisition, such as

changes in the specific scanner or acquisition parameters. The method is capable of generating samples for the normal tissues, including normal appearing white matter (NAWM), gray matter, and cerebrospinal fluid, as well as the pathological "tissue" enhancing lesion. The "samples" that it generates for the pathological "tissue" referred to as "lesion" are implicit from the above samples. The method is designed to deliver input to lesion finding and change detection methods. The output, however, is equally and fully applicable to a variety of other image processing methods, including registration, inhomogeneity correction, tessellation of surfaces for rendering or other purposes, segmentation for volume computation or other purposes, and the like. The method can be fully automated in that it does not require any user input and it is designed to be flexible enough that it does not require that image volumes supplied to it be acquired using particular hardware or acquisition parameters.

[0025] The sample point generation method, in contrast with more sophisticated atlas based segmentation methods, makes only very limited assumptions regarding anatomy, and is thus very flexible with respect to changes due to acquisition, disease, and interventions, such as resections. Use of simple but reliable knowledge makes the method very likely to provide highly accurate results. The present method, if used in a pipeline or spiral step prior to an atlas segmentation method, can also provide a means for an atlas based segmentation algorithm to understand some presented images more clearly, in spite of acquisition related variations. It allows the atlas-based algorithms to understand the images in terms of tissues such as gray matter, white matter, CSF, lesions, and enhancing lesions, instead of in terms of image intensities.

[0026] Another aspect of the invention is a method that identifies lesions, both T1-T2 abnormalities, and enhancing lesions, in the white matter of the brain, and to describe these lesions quantitatively. The lesions identified by this method may originate from primary or metastatic tumors, multiple sclerosis, or some other white matter disease process. The method can be fully automated and reproducible, and is highly robust against variations in acquisition parameters, the particulars of the scanner, and the like. From the original input volumes, for example, T1, T1-Post Gadolinium, T2, and the like, the method generates membership images for the classes, such as, lesion and enhancing lesion. On a voxel-by-voxel basis, these memberships vary in value from 0.0 to 1.0. As such, they may be straightforwardly used to compute measures of lesion load.

[0027] The lesion identification method is specifically designed to be most accurate in quantifying the memberships of voxels. To the greatest degree possible, the partial membership model is not intended to cater to a priori probabilities, but instead, the algorithm's model of partial-membership effects are derived from a physical model of the behavior of the imaging system. The intended purpose of the method is to quantify these partial memberships and partial characters in only two classes, for example, white matter lesion and enhancing lesion of the brain in magnetic resonance imaging studies. By focusing on this narrow but important topic, the algorithm is able to take advantage of a developed knowledge base.

[0028] The lesion identification method has been designed for the specific purpose of producing reproducible measures of lesion and enhancing lesion, with minimal or no user intervention. In contrast with previous algorithms, this method has been designed to be able to accurately quantify lesion and enhancing lesion membership even when no examples of frank lesion or enhancing lesion exist in the volume, for example, the algorithm performs supervised classification even when there are no examples. The method is relatively unaffected by changes in scanner and acquisition parameters.

[0029] Using the sample point generator, which identifies the white matter samples, the CSF samples, and the maximum enhancement automatically, the lesion detector has been shown to be highly reproducible, as well as automatic. The system provides an automatic and deterministic inhomogeneity correction, registration, and parenchyma mask generation. Although included in the detailed description below, these are in fact separate processes. When the processing pipeline uses deterministic sub-algorithms, the present change detection method yields consistent results every time it is run with the same volume data, regardless of who runs the application, where it is run, or how many times it is run.

[0030] Yet another aspect of the present invention is a method for comparing two MRI images of the head and identifying and characterizing changes that have occurred in white matter lesions that may be present. The method produces a color-coded change map, which is displayed superimposed upon the anatomical images. This color-coded change map identifies what has changed in the interval between the two MRI acquisitions, in what way, by how much, and where. The method aids in the accurate and simple identification of changes in serial brain MR images. The system provides a mechanism to ensure that all changes within serial MR images are observed and understood. The system provides a mechanism for defining

progression, for example, with regard to brain tumors, objectively. The system provides a mechanism for quantitatively comparing response to trial therapies across treatment groups. The system additionally provides a mechanism to promote reproducible diagnoses. The system may help to train new practitioners. It may serve as a model for how computers can liberate radiologists from doing mundane things sub-optimally, for example, measuring tumors with rulers on plastic film, which has been shown to be relatively ineffective. Instead a computer wades through mountains of analyses and presents the clinician only with refined, concise, and note-worthy data that has been tailored to their purpose. The method is a model of how computers can perform a large number of analyses and bring to a clinician's attention things they weren't looking for, but that are nonetheless important. The method additionally serves as a model of "layered analyses" that are impractical or impossible without computer assistance.

[0031] Another aspect of the present invention is a method for the detection of regions of signal, which uses a strong spatial connectedness criteria, and a variable threshold that weighs a region's spatial extent against the values of a given region's constituent voxels, to identify objects in a manner that possesses much higher sensitivity and specificity compared with simpler methods of separating signal from noise, such as thresholding. The method is able to identify objects made of up elements with intensities within the range of the background noise. It is able to identify such low-intensity objects of arbitrary shape, unlike methods that have used kernels of fixed size and shape. Exploration and identification of arbitrarily sized and shaped regions within a noisy background is potentially computationally intractable. The particular implementation described below makes the method tractable on mainstream computing hardware. As a component of a change detector system, this method has been applied to digital phantoms, and to serial MR scans of brain tumor patients, and it has shown very high accuracy, sensitivity, and specificity.

[0032] This invention is based in part on the realization that there are situations where domain-specific knowledge exists that can be leveraged to identify regions of signal lower than the noise floor. This knowledge could take a variety of forms. For example, in the case outlined in this application, it is known a priori that the three dimensional (3D) object that the image represents, is a real-world object, for example, a patient's brain. Within this brain, various types of objects or regions exist. An example of such an object or region could be a region of T2 abnormal white matter, or a region which is exhibiting increasing, that is, changing, T2

abnormality. In such cases, high intensity, either in a lesion membership image, or in a change in lesion membership image, is only one trait signifying a particular voxel's membership or lack of membership in the object or region of interest. This is because the object might possess only subtle intensity relative to the background. However, in many such cases, it is known *a priori* that objects, including subtle objects, frequently extend to encompass the space of multiple or many voxels, and are more or less continuous over the region which they cover. The extent of the region might serve as a second piece of knowledge, suggesting its membership in the region, while the continuousness of the region, that is, the fact that it is not thin and winding, and it is not full of holes, may serve as a third piece of knowledge. These three pieces of knowledge constitute the foundation for separating objects from background noise.

[0033] Various other features of the present invention will be made apparent from the following detailed description and the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0034] Fig. 1 is a block diagram of a typical prior art image comparison;

[0035] Fig. 2 is a block diagram of a more sophisticated prior art image comparison method;

[0036] Fig. 3 is a block diagram of an MRI system which employs the present invention;

[0037] Fig. 4 is a flow chart of the major elements in the preferred embodiment of a brain lesion change detector system;

[0038] Fig. 5 is a high level flow chart of an automated sample point finding procedure used in the procedure of Fig. 2;

[0039] Fig. 6 is a flow chart of the CSF finding procedure used in the procedure of Fig. 5;

[0040] Fig. 7 is a flow chart of the white matter and gray matter finding procedure of Fig. 5;

[0041] Fig. 8 is a flow chart setting forth the steps for of the maximum enhancement level determination procedure used in the procedure of Fig. 5;

[0042] Fig. 9 is a flow chart of the lesion identification process that forms part of the system of Fig. 4

[0043] Figs. 10-15 illustrates computation of membership in the class "enhancing lesion," where Fig. 12 shows the region of T1-T1-Post space considered to contain

enhancing white matter lesions (darkened), Fig. 13 shows a sample voxel that is mildly T1 hypointense, and which possesses moderate enhancement, Fig. 14 illustrates that in order to compute the membership of a trial voxel in the class "enhancing lesion," the voxel is first projected vertically onto the CSF-NAWM line, and Fig. 15 shows the vertical distance of the point from the "Max-Enh" line, and the vertical distance of the point from the CSF-NAWM line, are together used to compute the membership of the point in the class 'enhancing lesion';

[0044] Fig. 16 is a block diagram showing the change detector method in a system such as that of Fig. 4;

[0045] Fig. 17 is a block diagram of the change detector in the system of Fig. 16;

[0046] Fig. 18 is a screen capture of an exemplary output of the change detector;

[0047] Figs. 19-21 illustrate images and a histogram of the images to indicate the problem to be solved;

[0048] Fig. 22 is a flow chart of the object boundary identification method that forms part of Fig. 4;

[0049] Fig. 23 is a graph illustrative of the process in Fig. 22;

[0050] Fig. 24 is a pictorial representation of the process in Fig. 22;

[0051] Fig. 25 is a flow chart of a step in the process of Fig. 22 construction of the list of valid "snakes;"

[0052] Fig. 26 is a graphical representation of the steps taken during the process described with respect to Fig. 25;

[0053] Fig. 27 is a flow chart setting forth the steps in the procedure of Fig. 22 in which it is possible to use the list of snakes to produce a look-up table facilitating very rapid assessment of any neighborhood configuration and with this look-up table; and

[0054] Fig 28 shows a sample application of the lookup table to a neighborhood.

DESCRIPTION OF THE PREFERRED EMBODIMENT

[0055] Referring particularly to Fig. 3, the preferred embodiment of the invention is employed in an MRI system. The MRI system includes a workstation 10 having a display 12 and a keyboard 14. The workstation 10 includes a processor 16 that is a commercially available programmable machine running a commercially available operating system. The workstation 10 provides the operator interface that enables scan prescriptions to be entered into the MRI system.

[0056] The workstation 10 is coupled to four servers. Specifically, the workstation 10 includes a pulse sequence server 18; a data acquisition server 20; a data processing server 22, and a data store server 23. In the preferred embodiment the data store server 23 is performed by the workstation processor 16 and associated disc drive interface circuitry. The remaining three servers 18, 20 and 22 are performed by separate processors mounted in a single enclosure and interconnected using a 64-bit backplane bus. The pulse sequence server 18 employs a commercially available microprocessor and a commercially available quad communication controller. The data acquisition server 20 and data processing server 22 both employ the same commercially available microprocessor and the data processing server 22 further includes one or more array processors based on commercially available parallel vector processors.

[0057] The workstation 10 and each processor for the servers 18, 20 and 22 are connected to a serial communications network. This serial network conveys data that is downloaded to the servers 18, 20 and 22 from the workstation 10 and it conveys tag data that is communicated between the servers and between the workstation and the servers. In addition, a high speed data link is provided between the data processing server 22 and the workstation 10 in order to convey image data to the data store server 23.

[0058] The pulse sequence server 18 functions in response to program elements downloaded from the workstation 10 to operate a gradient system 24 and an RF system 26. Gradient waveforms necessary to perform the prescribed scan are produced and applied to the gradient system 24 that excites gradient coils in an assembly 28 to produce the magnetic field gradients G_x , G_y and G_z used for position encoding NMR signals. The gradient coil assembly 28 forms part of a magnet assembly 30 that includes a polarizing magnet 32 and a whole-body RF coil 34.

[0059] RF excitation waveforms are applied to the RF coil 34 by the RF system 26 to perform the prescribed magnetic resonance pulse sequence. Responsive NMR signals detected by the RF coil 34 are received by the RF system 26, amplified, demodulated, filtered and digitized under direction of commands produced by the pulse sequence server 18. The RF system 26 includes an RF transmitter for producing a wide variety of RF pulses used in MR pulse sequences. The RF transmitter is responsive to the scan prescription and direction from the pulse sequence server 18 to produce RF pulses of the desired frequency, phase and pulse

amplitude waveform. The generated RF pulses may be applied to the whole body RF coil 34 or to one or more local coils or coil arrays.

[0060] The RF system 26 also includes one or more RF receiver channels. Each RF receiver channel includes an RF amplifier that amplifies the NMR signal received by the coil to which it is connected and a quadrature detector that detects and digitizes the I and Q quadrature components of the received NMR signal. The magnitude of the received NMR signal may thus be determined at any sampled point by the square root of the sum of the squares of the I and Q components:

$$M = \sqrt{I^2 + Q^2},$$

and the phase of the received NMR signal may also be determined:

$$\phi = \tan^{-1} Q/I.$$

[0061] The pulse sequence server 18 also optionally receives patient data from a physiological acquisition controller 36. The controller 36 receives signals from a number of different sensors connected to the patient, such as ECG signals from electrodes or respiratory signals from a bellows. Such signals are typically used by the pulse sequence server 18 to synchronize, or "gate", the performance of the scan with the subject's respiration or heart beat.

[0062] The pulse sequence server 18 also connects to a scan room interface circuit 38 that receives signals from various sensors associated with the condition of the patient and the magnet system. It is also through the scan room interface circuit 38 that a patient positioning system 40 receives commands to move the patient to desired positions during the scan.

[0063] It should be apparent that the pulse sequence server 18 performs real-time control of MRI system elements during a scan. As a result, it is necessary that its hardware elements be operated with program instructions that are executed in a timely manner by run-time programs. The description components for a scan prescription are downloaded from the workstation 10 in the form of objects. The pulse sequence server 18 contains programs that receive these objects and converts them to objects that are employed by the run-time programs.

[0064] The digitized NMR signal samples produced by the RF system 26 are received by the data acquisition server 20. The data acquisition server 20 operates

in response to description components downloaded from the workstation 10 to receive the real-time NMR data and provide buffer storage such that no data is lost by data overrun. In some scans the data acquisition server 20 does little more than pass the acquired NMR data to the data processor server 22. However, in scans that require information derived from acquired NMR data to control the further performance of the scan, the data acquisition server 20 is programmed to produce such information and convey it to the pulse sequence server 18. For example, during prescans NMR data is acquired and used to calibrate the pulse sequence performed by the pulse sequence server 18. Also, navigator signals may be acquired during a scan and used to adjust RF or gradient system operating parameters or to control the view order in which k-space is sampled. And, the data acquisition server 20 may be employed to process NMR signals used to detect the arrival of contrast agent in an MRA scan. In all these examples the data acquisition server 20 acquires NMR data and processes it in real-time to produce information that is used to control the scan.

[0065] The data processing server 22 receives NMR data from the data acquisition server 20 and processes it in accordance with description components downloaded from the workstation 10. Such processing may include, for example: Fourier transformation of raw k-space NMR data to produce two or three-dimensional images; the application of filters to a reconstructed image; the performance of a backprojection image reconstruction of acquired NMR data; the calculation of functional MR images; the calculation of motion or flow images, etc.

[0066] Images reconstructed by the data processing server 22 are conveyed back to the workstation 10 where they are stored. Real-time images are stored in a data base memory cache (not shown) from which they may be output to operator display 12 or a display 42 that is located near the magnet assembly 30 for use by attending physicians. Batch mode images or selected real time images are stored in a host database on disc storage 44. When such images have been reconstructed and transferred to storage, the data processing server 22 notifies the data store server 23 on the workstation 10. The workstation 10 may be used by an operator to archive the images, produce films, or send the images via a network to other facilities.

[0067] Referring particularly to Fig. 4, the present invention employs the above-described MRI system to acquire a number of images from a subject's brain as indicated at process block 200. These are 3D images are referred to herein as

"volumes." The particular pulse sequences used will be listed below and suffice it to say that they are conventional acquisitions commonly employed in brain imaging.

[0068] Each of these volumes is first processed using a third-party inhomogeneity correction algorithm as indicated at process block 202. For example, N3 may be used, but other solutions may be used, as well. Next, the volumes are spatially registered to one another as indicated at process 204. In accordance with one aspect of the invention a proprietary solution may be used that is a hybrid of other third-party registration algorithms. From these registered volumes, a parenchyma mask is automatically generated by, for example, another third party software application as indicated at process 206. In accordance with one aspect of the invention, "Brain Extraction Tool" (BET) or any other appropriate tool may be used.

[0069] As indicated at process block 259, the tissue sample generator is then employed to produce from the acquired image volumes sample pixel values for white matter, gray matter, and CSF. In addition, it finds the maximum enhancement level produced in the images by an administered contrast agent. The structure and operation of the tissue sample generator is described in detail below.

[0070] As indicated in Fig. 4, at process block 257, a lesion identification process is then performed. This employs the acquired image volumes and the calculated tissue samples to locate and quantify any abnormal tissues. The structure and operation of the lesion identification process is described in detail below.

[0071] As indicated at process block 260, the next step is to detect changes that have occurred in identified lesions since the last examination. The process 260 accomplishes this by comparing the current identified lesions with lesions identified in previous examinations that have been stored for this purpose. The operation and structure of the lesion change detector is described in detail below.

[0072] And finally, an object boundary identification process indicated at process block 258 may be employed to refine the boundaries of identified lesions or refine indicated changes in lesions. This is a noise reduction filter that is described in more detail below.

[0073] The goal is to provide an automated image analysis and the process performed by this system may be viewed as a spiral of knowledge. That is, it starts with very simple but very reliable knowledge at the center of the spiral then, continuing around the spiral, each loop of the spiral builds upon knowledge generated by the prior loops of the spiral. Each loop of the spiral embodies more and more sophisticated processing, both using and producing progressively more

sophisticated knowledge. In the inner-most loop, there might be automated sample point systems of the type described. In a middle loop, automated segmentation results are provided that divide the brain into white matter, gray matter, CSF, lesion, enhancing lesion, and the like. Then in a yet further loop, volumetric and other measures of particular structures based upon the segmentation results are provided. In loops even further towards the outside, suspected conditions, profiles, risk factors for the particular patient, based upon what has been measured in the images are provided.

[0074] Tissue Sample Generator

[0075] Fig. 5 shows a high level flow of execution for the automated sample point finding method. As shown in resource block 300, the method uses T1, T1-Post Gd, T2, and PD images. If the T1-Post Gd sequence has been acquired with magnetization transfer suppression (MTS), then a T1 sequence without contrast that has been acquired with magnetization transfer suppression is also supplied, as also reflected in resource block 300. Additionally, resource block 300 also includes a parenchyma mask. As shown in process blocks 302 and 304, the method first finds cerebrospinal fluid (CSF) samples and uses these samples in the process for finding white matter and gray matter samples. Finally, as shown in process block 306, the system separately finds the maximum enhancement level. As shown in results block 308, the product of the method is the multi-spectral, that is, for all of the supplied pulse sequences, mean and standard deviation of white matter, gray matter, and CSF. Furthermore, the method additionally finds the maximum enhancement level for white matter, which, by definition is uni-spectral, since enhancing tissue may present with a range of intensities in pulse sequences other than T1-Post.

[0076] Fig. 6 shows the organizational flow of the CSF finding sub-routine of the automated sample points generator. The method uses as input a T1 weighted MR image, a T2 weighted MR image, and a parenchyma mask, as reflected in resource block 400. The method uses three pieces of simple but reliable knowledge to find CSF samples. First, it recognizes that CSF is one of the darkest things in a T1 weighted image and one of the brightest things in a T2 weighted image. Second it recognizes that there is a lot of CSF in an MR image of a patient's head. Finally, it recognizes that CSF is found in large continuous spatially connected regions in a head MRI. As long as these three criteria are met, the method will be able to identify samples of CSF, regardless of the patient, the pathology, whether or not a resection

has been performed, what scanner the images have been acquired with, or what particular acquisition parameters have been used.

[0077] The CSF finding method begins with two branches. One branch 402 operates on the T2 volume. The other branch 404 operates on the T1 volume. The actual operation of these two branches 402, 404 is quite similar. In both branches 402, 404, the algorithm constructs a non-air mask; however, the first mask uses the T2 weighted image, as indicated at process block 406, and the second mask uses the T1 weighted image, as indicated at process block 408. Using a known process, which essentially finds an intensity threshold that separates air from biological tissues, that is, the patient's head. The process is based on the assumption that there will be a lot of air and the imaging intensity of this air will be quite low. Both branches of the CSF finding procedure then constructs a histogram of the intensities of the non-air tissues with, for example, 200 bins, as indicated at process block 410, 412.

[0078] Thereafter, both branches then walk the histogram. In the first branch 402, which relates to the T2 weighted volume, the process begins with the highest intensity and walks the histogram towards progressively decreasing intensities, until it finds a bin whose height exceeds the number of voxels in the T2 non-air mask, computed above, divided by 700, as indicated at process block 414. This step embodies the knowledge that in a head MR image, CSF voxels are the largest group of very bright voxels in this volume. In the first branch of the CSF finding sub-routine 402, the process then uses this intensity level as a threshold at process block 416 to construct a mask by finding all of the voxels in the T2 weighted volume with intensities exceeding this level.

[0079] In the histogram walking step of the second branch 404, the method is by contrast attempting to find a large number of voxels with low intensity. The process continues by walking with the tallest bin in the image, and walks to the left / towards lower intensity, until it reaches a bin whose height is less than 0.015 times the number of voxels in the T1 non-air mask, as indicated at process block 418. The process continues by searching for the bin between this bin, and the lowest intensity bin, such that using the intensity of the bin as a threshold, as indicated at process block 420. For example, the process can construct a mask of all voxels in the T1 volume with an intensity below this intensity. This results in the largest object in the image to be as close as possible to 600 voxels in extent. Having computed this intensity, the second of the CSF finding process 404 then creates a mask containing

only the voxels whose intensities in the T1 weighted image are less than this threshold, as indicated at process block 422. The process shown in the second branch also computes the intersection of this mask, the non-air mask, and the parenchyma mask to construct a new mask of CSF candidates. The logical intersection of the mask from the first branch 402 and the mask from the second branch 404 is then computed at process block 424. From this unified mask, the process eliminates all objects possessing fewer than 300 voxels at process block 426. From the remaining objects, at process block 428, the method then identifies the object that minimizes the following equation, which essentially attempts to identify the group of voxels in the mask that are brightest in T2 and darkest in T1, while sensibly balancing these two objectives with regard to one another.

$$\text{Cost} = \frac{T1\text{ObjectMean}}{\text{MaxValueT1Parenchyma} + (\text{MaxValueT2Parenchyma} - T2\text{ObjectMean})d / \text{MaxValueT2parenchyma}} \quad (1)$$

[0080] Fig. 7 shows the organizational flow by which the method identifies samples of normal appearing white matter and gray matter. As in the CSF finding method described above, the NAWM/gray matter finding method uses simple but reliable knowledge in order to identify tissue samples. In this case, the knowledge includes the fact that both normal appearing white matter and gray matter are present in a head MR image in large spatially contiguous regions and both normal appearing white matter and gray matter represent a significant fraction of the patient's total head. In addition the knowledge includes the fact that, in a T1 weighted image, both normal appearing white matter and gray matter are substantially brighter than CSF, for example, the intensities distributions of normal appearing white matter and gray matter do not overlap with the intensity distribution of CSF. Similarly, in a T2 weighted image, both normal appearing white matter and gray matter are substantially darker than CSF and, again, their distributions do not overlap with the intensity distribution of CSF. The normal appearing white matter/gray matter finding sub-routine uses as input T1, T1-Post, T2, and PD sequences, as reflected by resource block 500. Additionally, if the supplied T1-Post Gd image has been acquired with magnetization transfer suppression (MTS) turned on, then an additional T1 volume (without contrast) which also has magnetization transfer suppression turned on should be supplied. The method also uses a

parenchyma mask, and the mask identifying the CSF samples created according to the process shown in Fig. 6.

[0081] The NAWM/gray matter sample finding method first computes the mean and standard deviation of CSF in all pulse sequences at process block 502, as well as the min and max intensities of CSF at process block 504. It additionally computes a non-air mask from the PD volume, at process block 506, using the same process referred to above. At process block 508, the process zeros all voxels in this mask for which either the parenchyma mask is not set, or for which the intensity in the T1 volume is within three standard deviations of the CSF mean in T1, or for which the intensity in the T2 volume is within three standard deviations of the CSF mean in T2. The latter two stipulations roughly eliminate all the CSF voxels from consideration. The process then explores a variety of trial intensities at process block 510, in an attempt to find large spatially contiguous regions with very similar intensities. Specifically, for each center intensity "Trial_{PD}," it identifies all voxels within one standard deviation of the trial intensity, where the CSF standard deviation in the PD volume is used to define the appropriate level of noise. It then enforces a strong spatial connectivity constraint, as in the noise reduction method described below, to split all resulting regions into spatially distinct regions. From the collection of identified regions, the process discards all regions smaller than 600 voxels. After these steps have been completed, for each Trial_{PD} intensity, the process records the total number of voxels. The Trial_{PD}, number voxels pairs are then interpreted as a histogram, for example, the intensity Trial_{PD} represents the X-axis of the histogram, and the number of voxels identified at that intensity is used as the bin height. The process identifies the tallest bin and uses its intensity as the gray matter intensity, as indicated at process block 520. The process then walks the histogram at process block 522, starting from the lowest intensity and progressing towards the highest intensity, until it finds a bin whose height is greater than 0.08 times the height of the tallest bin in the histogram. The intensity of this bin is used as the normal appearing white matter intensity.

[0082] The method generates white and gray matter maps at process block 524 using the two intensity levels determined above. For example, the intensity possessing the tallest bin in the histogram, and the lowest intensity for which the bin height exceeds 0.08 times the tallest bin, including voxels which fall within one standard deviation of those intensities, where the noise level used is the same noise level as that of CSF in PD, which meet the spatial connectivity constraint, and which

possess at least 600 voxels per strongly spatially connected clump. The process generates a histogram for each tissue, for example, white matter and gray matter, and from each pulse sequence, and identifies the peak in each histogram, as indicated at process block 526. These peak intensities are used as the final center intensities for normal appearing white matter and gray matter in each pulse sequence and the standard deviation of the noise of the original CSF samples is used as the noise level in all cases, as shown at process block 528, to yield the desired results reflected at resource block 530.

[0083] Referring now to Fig. 8, the organization flow of the method that determines maximum enhancement level is shown. Its operation is similar to that of the other automated sample point finding methods, although, in some ways moderately simpler in that, as with the other methods, this method leverages simple but reliable a priori knowledge. In this case, the knowledge that serves as a foundation for the determination of maximum enhancement level includes the fact that there are a relatively large number of voxels which enhance in a T1-Post image and that enhancing regions are, by definition, the brightest voxels in a T1-Post image. It should be noted that these enhancing regions, include extra-cranial soft tissues, dura, if it enhances, blood vessels, and the like. These regions, even though not neurological white matter themselves, possess intensities that are representative of the maximum degree to which white matter will enhance, as the blood brain barrier breaks down. As indicated at resource block 600, the method uses the intensity of these normal tissues in T1-Post images as the basis for computing the maximum enhancement level, which means that this sub-routine can reliably compute maximum enhancement level, even when no frankly enhancing white matter tissue exists in the image.

[0084] The method for determining maximum enhancement intensity requires only a T1-Post volume and it is not important from the perspective of this method whether or not the volume was acquired with magnetization transfer suppression enabled. As indicated at process block 602, the method involves computing a non-air mask from the T1-Post volume. This is done using the same process as described above. Next, the number of voxels in the non-air mask is computed at process block 604. Following this, at process block 606, a histogram is constructed from the intensities of the non-air voxels. Finally, the histogram is scanned from highest intensity to lowest intensity until it locates a bin whose height equals or exceeds the number of non-air voxels in the image divided by 1044, as indicated at

process block 608. The intensity of this bin is used as the maximum enhancement intensity, as indicated at resource block 610.

[0085] Accordingly, the present invention takes an entirely different approach to prior art analyzing systems. Specifically, the approach described here uses the knowledge that white matter lesions, in the extreme, always come to resemble CSF, and thus CSF, which is always present in MR images of the brain, serves as a useful and consistent end-point for lesion. Likewise, the present method is able to determine a useful measure of maximum enhancement from the images, in a way that is very consistent and reliable, even when there are no frank enhancing lesions, because it is able to base the measure of maximum enhancement on the enhancement of normal tissues, which are always present. The present method also abandons the idea of enhancing tissue as a unique tissue with its own multispectral imaging signature. This enhancing tissue may be present in a variety of ways across pulse sequences other than the T1-Post Gadolinium (Gd) image. For example, fully enhancing tissue may be very abnormal in T2, or it may be only slightly abnormal in T2, and the like. Likewise, enhancing tissue may be nearly normal appearing in T1, or it may be very T1 hypointense. Thus, in contrast with previous approaches, enhancement is considered to be special with respect to other pure tissues, in that its description is uni-spectral. Another difference between this method and prior methods is that the present method does not attempt to characterize the imaging characteristics of necrosis. Although necrosis is real, its actual cytological manifestations are too variable to be characterized by anything like a single unique centroid in imaging intensity feature space, and therefore it has been omitted.

[0086] The automated sample points generator provides input to the lesion finder and change detector algorithms described below. Certainly there are many other procedures that either use or could benefit from automatic sample point generation. One example is classification algorithms. Other less obvious examples include registration, inhomogeneity correction, anatomical atlas based segmentation, finite element model development, and the like. Particularly since the present method uses *a priori* knowledge that is as limited as possible, it should find application very early in any processing pipeline, and thus can serve as an important source of knowledge for other subsequent, higher-level stages of processing.

[0087] Lesion Identification

[0088] Referring now to Fig. 9, the high-level architecture of the lesion identification and quantification process is shown. The process resembles a

classification algorithm in many respects; however, it has several important differences. The first and most important distinguishing feature is that it is not general-purpose. Rather, it is designed specifically to accurately quantify membership in white matter lesions of the brain in magnetic resonance imaging (MRI) data. Focusing on this domain in particular, allows a level of performance that would be unachievable using general-purpose algorithms. Another feature that distinguishes this method from standard classification algorithms, is that it is intentionally focused on accurately quantifying partial membership and partial character mixtures, instead of identifying voxels which contain pure tissues.

[0089] As shown in resource block 700, the procedure uses as input T1, T1 Post Gd, and T2 volumes. Additionally, if the T1-Post Gd volume which is supplied has been acquired with magnetization transfer suppression turned on, then a T1 volume acquired with magnetization transfer suppression but without Gd should also be supplied, labeled 'MTS' in Fig. 9. Finally, a proton density (PD) volume may be supplied, which facilitates the fully automated tissue sample generator as described above in which the system automatically determines the normal-appearing white matter samples, CSF samples and maximum enhancement level instead of having the user supply these samples. Using the samples, the method performs inhomogeneity correction at process block 702 and registration at process block 704, in a manner described above. Also, in a manner described above, a parenchyma mask may be automatically generated at block 706 or a user may supply NAWM samples from anterior corpus callosum at process block 708. In any case, at process block 710, samples of CSF, if applicable, NAWM, and maximum enhancement are determined. Then, at process block 712 lesion and enhancing lesion memberships are computed on a voxel-by-voxel basis by a process that is described further, below. The process finally applies a noise reduction sub-algorithm at process block 714, which also described below, to identify which regions actually exhibit abnormal character, as opposed to exhibiting abnormal character only due to noise, and yields the desired output 716.

[0090] Referring now to Figs. 10 and 11, the method used to compute T1-T2 abnormality can be explained. This step is first premised upon the observation that T1, for example, T1-hypointensity and T2, for example, T2 hyperintensity are exhibited in unison. Second, it is premised on the fact that T1-T2 abnormal white matter transitions in such a way that the more abnormal it becomes, the more its T1-T2 profile comes to resemble the T1-T2 profile of CSF. In some cases, the tissue

first acquires T2 hyperintensity and only after it has acquired some degree of T2 hyperintensity does it begin to acquire T1 hypointensity and T2 hyperintensity concurrently. In other cases, the tissue acquires T1 hypointensity and T2 hyperintensity at the same time, from the first stages of abnormality acquisition. In all cases, white matter in all stages of normality and abnormality may be seen to exhibit T1-T2 profiles within the gray polygon shown in Fig. 10. In this figure, the location of the NAWM and CSF centroids are indicated by small circles. Ideally, the T1-T2 profile of abnormal white matter would lie within the portion of this polygon to the left of the NAWM centroid; however, blood can elevate the T1 value of white matter voxels in some cases and cause them to lie within the region of the polygon to the right of the NAWM centroid. The method first identifies the voxels which lie within the polygon. The method then computes the Euclidean distance of each such voxel in T1-T2 feature space from the NAWM and CSF centroids, as shown in Fig. 11. Finally, the method uses these distances to compute the membership of the voxels in the class "lesion" using the following equation.

$$\text{membership}_{\text{lesion}} = \frac{\text{Dist}_{\text{PIToNAWM}}}{\text{Dist}_{\text{PIToCSF}} + \text{Dist}_{\text{PIToNAWM}}} \quad (2)$$

[0091] This method uses the automatically determined sample points to normalize the intensities, which provides a high degree of robustness against changes in acquisition parameters, scanner, and the like. The memberships computed are therefore quantitative, and reproducible. The method makes use of specific knowledge regarding white matter lesions, such as that T1 and T2 intensity signatures of white matter lesions act in unison. More specifically, as a region of white matter initially becomes abnormal, it exhibits T2 hyperintensity. T1 hypointensity may or may not accompany this initial T2 hyperintensity. As the T2 hyperintensity accumulates, T1 hypointensity begins to accumulate more rapidly. As mentioned, above, at the extreme limit of abnormality, white matter comes to resemble CSF in terms of its T1-T2 intensity profile. Furthermore, in the case of partial voluming, that is, the boundary between two regions of pure tissue transects a given voxel, the fractional membership in each of the tissues is equal to the fractional position along the line connecting the centroids of the two tissues: CSF and NAWM, in T1-T2 feature space. In the case of partial character, there is some deviation to the right of the line connecting the two centroids, although the feature space

fractional distance approach still provides a useful methodology for computing membership.

[0092] Referring now to Fig. 12, the method used to quantify membership in the class, "enhancing lesion," is demonstrated. First, the method uses the knowledge that T1, and T1-Post are the same pulse sequence, although they may possess different acquisition parameters, with the principal difference between the two images being that prior to acquisition of the T1-Post sequence, an intravenous contrast medium is administered, which causes regions of blood brain barrier (BBB) breakdown to exhibit heightened image intensity in the T1-Post image. The intensity of non-enhancing regions, that is, regions where the BBB is not in the process of breaking down, will therefore always lie along the line crossing the CSF and NAWM centroids in T1-T1-Post feature space. On the other hand, the intensity of voxels representing regions where the BBB is in a state of breakdown should lie above this line in T1-T1-Post feature space. In Fig. 12, the CSF and NAWM centroids are shown as small circles, and the line crossing these centroids is additionally shown. Furthermore, there is a maximal intensity which may be observed in the T1-Post sequence due to the presence of contrast medium within the space represented by a voxel. Within the range of contrast concentrations likely to be seen in enhancing tissue, the intensity in the T1-Post sequence of some voxel representing white matter is expected to lie either on the line crossing the CSF and NAWM centroids, for no enhancement. Above the line but below the horizontal "Max-Enh" line shown in the Fig. 12, for some degree of enhancement, or along the "Max-Enh" line, for maximal enhancement.

[0093] The degree of enhancement is quantified as if it were a linear effect. That is, the fractional distance along the line connecting the point's vertical projection on the CSF-NAWM line, and the point's vertical projection on the "Max-Enh" line, provide that voxel's membership in the class, "enhancing lesion." This process is shown graphically in Figs. 12-15. In Fig. 13 a trial point is shown by a small circle in the interior of the gray region. In order to quantify the degree of enhancement exhibited by the voxel, the voxel is first projected vertically on the CSF-NAWM line, as shown in Fig. 14. This provides the intensity in the T1-Post sequence that would be observed if the voxel exhibited no enhancement. The fractional distance of the point in question along the vertical line connecting this projection and the "Max-Enh" line is shown graphically in Fig. 15. The vertical distance to the "Max-Enh" line, and

the CSF-NAWM line are used to compute the voxel's membership in the enhancing lesion class using the following equation.

$$\text{membership}_{\text{EnhancingLesion}} = \frac{\text{Dist}_{\text{PtTo Proj}}}{\text{Dist}_{\text{PtToMaxEnh}} + \text{Dist}_{\text{PtTo Proj}}} \quad (3)$$

[0094] It should be observed that the trial voxel shown in Fig. 12-15 exhibits some degree of T1 hypointensity. This is clear from the fact that the trial point lies to the left of the NAWM centroid in the figures. However, because the method uses the T1 sequence as a basis against which the T1-Post intensity is judged, this hypointensity has no effect on the computation of enhancing lesion membership of that voxel. It should be noted, furthermore, that this normalization method works for elevated T1 intensities, as well. Blood in the region represented by a voxel does not have any effect upon the computed membership of the voxel in the class "enhancing lesion," for the same reason. It should further be noted that if the T1-Post Gd image were acquired with magnetization transfer suppression enabled, then the T1 sequence on the X-axis of Figs. 12-15 would be substituted with a T1 sequence also acquired with magnetization transfer suppression enabled. In this way, mildly bright regions, due to T2 lesions and not due to enhancement, would still lie along the line connecting the CSF NAWM line, but above and to the right of the NAWM centroid, and would be correctly quantified as not possessing any enhancement.

[0095] Lesion Change Detection

[0096] The lesion change detection process includes a pipeline of steps shown in Fig. 16. Most of the elements in this figure include pre-processing steps described above. As indicated at process block 800, the process begins with the acquisition of images in a manner consistent with the objective, in this case change detection. This means, for example, that intravenous (IV) contrast should be administered in both acquisitions at the same time relative to image acquisition, or at a time during the image acquisition such that any differences between the acquisitions do not impact the appearance of the contrast in the images. Other acquisition-related issues include: reducing the slice thickness, in particular, to limit the number of tissues which might be included in individual voxels; eliminating inter-slice gaps because the gaps correspond to regions of the brain where there is no information from which to compute changes, and more importantly, the presence of inter-slice gaps make it impossible to correct for registration errors in the through-slice

direction, performing three dimensional (3D) acquisition; acquiring the images using an equivalent scanner, scanner software, pulse sequences, and acquisition parameters; correcting for decay of gradient coils; correcting for inhomogeneities at the time of acquisition to the greatest degree possible; and acquiring the images as close to spatially registered as possible, since frequently registration algorithms introduce artifacts in proportion with the degree of difference between the images being registered.

[0097] After the images have been acquired, intensity inhomogeneities are generally still present in the images. As described above, the present system attempts to resolve these using, for example, the N3 software package at process block 802. The images are then spatially registered to establish a common spatial frame-work at process block 804. The system uses a hybrid algorithm consisting of off-the-shelf registration algorithms developed by other groups, but assembled by a programmer in our group. At process block 806, a parenchyma mask is then generated using a tool called "Brain Extraction Tool" (BET). Following this, the actual comparison between the acquisitions is performed at process block 808. Finally, images are presented as a color-coded change map superimposed upon the patient's anatomical images at process block 810. This change map shows what is changing, where, in what way by the color, and by how much by the intensity of the color. Additional quantitative summaries of the changes that are present are generated. These may used to provide a global summary of the changes that are taking place. Additionally, in the past, relatively simple heuristics and mathematical rules have been used to process the output of the change detection algorithm, in order to differentiate between stable and progressing cases. Such approaches have shown themselves to be extremely effective and useful.

[0098] Referring now to Fig. 17, the lesion change detection method described above includes three primary steps. First, an automated sample points algorithm, which automatically defines samples of normal-appearing white matter (NAWM), cerebrospinal fluid (CSF), and gray matter is performed at process block 812 and will be described in detail below. Additionally, the automated sample points method automatically generates a mathematical description of the lesion and enhancing lesion imaging properties in the particular images. The method makes this possible even when there are no frank examples of lesion or enhancing lesion in the images. Next, at process block 814, the system measures of change are computed on a voxel-by-voxel basis. These are expressed in normalized Δ membership form.

Finally, at process block 816, a noise reduction step is performed as will be described in detail below. Due to noise and other factors, virtually all voxels will demonstrate changes of some sort. Only changes due to the underlying disease process are of interest, however, and therefore the purpose of this step is to separate disease related changes from non disease related changes. In order to do this, the method automatically identifies strongly spatially connected regions that are changing in the same way, and then applies a kind of test of "significance," such that changes small in magnitude but of large spatial extent may be identified as significant, and likewise changes small in spatial extent but large in magnitude may also be correctly identified as real.

[0099] In order to facilitate real-world application of the steps shown in Fig. 16, an application was created that unifies all of the processing steps, into a single very user-friendly system, and that additionally facilitates interaction with the output of the change detection algorithm using the output of the change detector algorithms and some additional features. A screen capture from the output screen of this application is shown in Fig. 18. In addition to providing a user-friendly mechanism for the user to select images upon which to perform change detection, and actually running each of the steps shown in Fig. 16, the change detection application facilitates validation of the output of each of the pre-processing steps shown in Fig. 16. For example, the registration algorithm we use is quite reliable, although in some instances it fails and must be re-run. The change detection application, therefore, provides a mechanism for the user to inspect the output of the registration algorithm, and allows the user to re-run the registration algorithm in the rare event that the results are unsatisfactory. Correspondingly, the change detection application allows the user to inspect the output of the automatic parenchyma mask generation program, and to edit the mask as appropriate. The output screen of the change detector application, an example of which is shown in Fig. 18, provides two modes of color change map overlay, which may be turned on and off. It provides linked cursors, so that the user may identify spatial correspondence between pulse sequences and the color change map and it provides a "flicker mode," so that the user may view a rapid alternation between baseline and follow-up examinations.

[00100] There are a number of important differences between this method and prior methods. First, prior automated sample point algorithms are only "maximally" automated, requiring the manual provision of samples of normal appearing white matter, while the current algorithm is automated, not requiring that any samples be

supplied by the user. This is a very important point, considering the time involved in defining samples, the variability that manual definition of these samples introduces, and the training issues involved. Second, prior algorithms often require the use of a specific imaging protocol. Generally, requiring that images be acquired with a very specific protocol is both unnecessarily restrictive, and unrealistic. The present method is extremely flexible with regard to acquisition parameters. Third, prior algorithms are only applicable to brain tumors, while the present method is applicable to multiple sclerosis, metastatic brain tumors from non-brain primaries, and other forms of white matter pathology. Fourth, prior algorithms use a fluid attenuation by inversion recovery (FLAIR) sequence whereas the current method uses T2, instead. Use of T2 instead of FLAIR makes the algorithm more sensitive, because T2 is monotonic with regard to lesion development. Generally speaking, the more pathological a region is, the higher its intensity in T2, which contrasts with FLAIR, where as a region of white matter becomes more abnormal, it first increases in intensity, but then subsequently decreases, as it acquires a more fluidic character. FLAIR additionally suffers frequently from fluid suppression artifact, a problem that T2 is not susceptible to since it does not suppress fluid. Generally speaking, the use of T2 in lieu of FLAIR makes the definition of lesion imaging characteristics much more consistent, and thus makes the algorithm much more robust, that is, sensitive, reliable, consistent, and flexible with regard to the particulars of specific image acquisitions. Fifth, the present method uses a very different model of lesion and enhancing lesion imaging characteristics, which is more "correct." The approach taken by the present method is thus more sensitive and accurate.

[00101] Object Boundary Identification

[00102] This section describes a noise reduction method used to identify spatially cohesive regions of signal embedded in a zero-mean field of noise. This essentially operates in two parts. First, the method identifies strongly spatially connected regions with values, in this case, image intensities, that are non-zero and which are on the same side of zero, that is, all positive or all negative. Second, it compares the mean value of each identified region with a variable threshold, to determine if the region is "significantly" different from zero. In this way, very small spatially connected regions may be detected if they contain sufficiently large voxel intensities, and likewise regions with very small intensities, even below the noise floor, may be identified if they cover a sufficiently large spatial extent. This presents a huge advantage over simple thresholding methods, which are not applicable to detecting

regions of signal below the noise floor, regardless of the region they cover. This is illustrated in Figs. 19-21. Specifically, Fig. 19 shows a square object surrounded by a background field of zero-mean noise. In this example, the CNR between the object and its background is very low. This kind of problem is prevalent in many contexts. As shown in Fig. 20, what is desired is to identify which voxels correspond to the object, and to eliminate the voxels corresponding to the noisy background, as shown. It should be noted that the image in this part is the gold standard, and is shown for explanatory purposes. Fig. 21 is a histogram of the intensities. When the CNR is so low that the histogram of the object and the histogram of the background overlap extensively, it is not possible to use thresholding to separate the two. This is a major problem in cases where very high specificity is required, such as lesion change detection, because in order to use thresholding to completely eliminate regions of zero-mean noise, visually obvious regions of signal would have to be eliminated as well.

[00103] In magnetic resonance imaging, the problem of partial-voluming makes this problem even more acute because “objects” are completely surrounded by a layer of partial-volumed voxels, which resemble the background even more than do the interior object voxels. These edge voxels are important to detect, just as are the interior voxels, particularly when the objects are very small and thus the edge voxels make up a great proportion of the total object. The purpose of the method described in this section, is to recognize objects such as the one shown, not just by their intensity but also by their spatial connectivity and extent.

[00104] The method is designed to detect and separate “object” voxels, from voxels containing zero-mean noise. That is, it attempts to identify regions that contain voxels that as a spatially-connected clump are “significantly” different from zero, in light of the level of noise present. An “on” region would therefore contain non-zero voxels all possessing the same sign in the volume being investigated. The method is not limited to unsigned data; however, the method may be applied to signed data, by searching for spatially connected regions twice: once for spatially-connected regions possessing positive intensities, and once for spatially-connected regions possessing negative intensities. The method operates in three dimensions, that is, it is able to locate three dimensional spatially contiguous regions, which improves its sensitivity compared with two dimensional algorithms, as regions of interest in tomographic medical imaging are typically three dimensional. Unlike a standard region-growing algorithm, the new method uses a strong spatial

connectivity constraint. In particular, to be considered connected to a region, voxels must have a number of nearest neighbors that also appear to be a part of the object, and the spatial relationship of these nearest neighbors with regard to one another is very particular. This strong spatial connectivity constraint prevents the method from identifying long winding regions in high background noise fields. The spatial connectivity constraint is implemented in the form of a spatial configuration lookup table, which is generated once, in advance. Without this look-up table, the process of enforcing strong spatial constraints would possess a great risk of being intractable. Although the method was designed with medical imaging, and particularly lesion change detection, in mind, it would be equally applicable to a large variety of fields. For example, the method would be useful within medical imaging including fMRI, classification, and others, and outside of medical imaging including remote sensing, astronomy, geology, and many others. The method is equally applicable to a large variety of source data types, including MRI, CT, radar, visible light, and the like.

[00105] Referring now to Fig. 22, the basic high level architecture of the method is shown. First, at process block 900, the operational data structures are either built, in the case of the simple structures, or loaded, in cases where the complexity of the structure is too great to be constructed at each instantiation, such as the case of the look-up table whose construction is described with respect to Figs 28 and 29. Next, the image volume that is to be explored for regions of signal is examined at process block 902. If regions with positive intensities are to be identified, then all voxels with positive intensities are loaded into a doubly linked list, which is subsequently sorted according to intensity.

[00106] A control loop is then commenced beginning at process block 904. At each instance of this loop, the next largest voxel is selected from the doubly linked list. At process block 906, the selected voxel coordinate is assigned a unique number, which will be referred to as an "intermediate index." The details and purpose of this intermediate index are discussed below. The intermediate index is recorded at the coordinates of the original voxel, but in a separate volume. A new clump info descriptor is created for the voxel, and appropriate data structures are updated. The term 'clump' is being used, here, to describe a strongly spatially connected region of signal. Now, at process block 908, the neighborhood of the new voxel is then examined, in order to determine whether the new voxel is sufficiently strongly connected, for example, sufficiently surrounded, by its neighbors to warrant

merging into a larger clump with them. This is the strong spatial connectivity requirement alluded to, above. The specific details of this constraint are relatively complicated and are extremely central to the functioning of this method. They are described in detail below.

[00107] At process blocks 910 and 912, the voxel and its clump, as it comes to belong to a clump, are merged with neighboring voxels and their clumps, if the neighboring voxels belong to clumps, as allowed by the strong spatial connectivity constraint. Furthermore, with the addition of each new voxel, the neighbors of the new voxel are re-tested for indicated merges from *their* perspective, since the addition of the new voxel may have been all that was needed to satisfy *their* spatial connectivity constraint. Note that each time a merge is performed, if the contributing clumps have already been judged to be significant, then the collective clump is also judged to be significant. At process block 912, if the new clump has not inherited significance from one of the clumps that contributed to its creation, the new collective clump is tested for significance, according to the function shown in Fig. 23. It should be underscored that once a clump has met the criteria of significance, it confers significance on any clump to which it is subsequently merged. This means that a small but significant region surrounded by a penumbra of very low intensity will be collectively significant, even if the region as a whole does not meet the significance criteria shown in Fig. 23. After this, process block 916 is reached and the main loop returns to its beginning, and selects the next largest voxel from the list, and the process is repeated.

[00108] The purpose of the data structure shown in Fig. 24 is to facilitate very efficient merging of clumps, which is a task performed a very large number of times during the execution of the main loop in Fig. 22. In addition to keeping track of statistics for each clump, for example, number of voxels, mean of the voxel values, and whether or not the clump has been judged to be significant, the structure is fundamentally circular. That is, using this structure, it can be determined very rapidly which clump a given voxel belongs to, and it can be likewise determined very rapidly which voxels belong to a clump. Furthermore, using this structure, the process of reassigning a list, potentially a very large list, of voxels from one clump to another is extremely rapid. The purpose of the volume of intermediate indices and intermediate index array, as opposed to simply using a volume of pointers, is reduced memory utilization. A 512x512x48 volume has over 12.5 million voxels. A 512x512x144 volume has almost 38 million voxels. In any given volume to be inspected, it is likely

that only a fraction of these will have “on” values. Based on this fact, and the likely difference in memory space required to store an index versus a pointer, the given approach was adopted. If memory utilization were not an issue, however, the concept of intermediate index could be abandoned, and the volume of intermediate indices and intermediate index array could be replaced with a volume of pointers to clump info descriptors.

[00109] The following section describes the construction of the operational data structures which are used to assess neighborhoods for the need for merging. Fig. 25 is a flow chart setting forth the steps of the process of constructing “snakes” that are named simply for the fact that they wind back and forth within a 3x3x3 neighborhood. The process starts at process block 1000 by creating a binary number, for example, having a length of 26. At decision block 1010, a check is performed to see if the snake length includes, for example 8 “on” elements. That is, the process confirms that the putative snake is of the correct length of 8. If not, at decision block 1011, the process increments the binary number by adding 1 to the least significant bit and a check is made to determine if binary number has reached overflow. If not, the process returns to process block 1010. A check is then performed to determine if the “on” neighbors are spatially continuous at decision block 1012. That is, the process confirms \that it is possible to traverse all 8 elements via full-face connections. If not, the process returns to decision block 1011. Once it is determined that the “on” neighbors are spatially continuous such that it is possible to traverse the entire snake vial full-face connections, a check is made to determine whether the snake is sufficiently compact at decision block 1014. Specifically, the process confirms that the number of full-face adjacencies within the snake is greater than one less than the length of the snake, which forces putative snakes to wind back on themselves at least once, and prevents purely linear snakes. If not, the process reverts to decision block 1011. On the other hand, if the number of intra-snake, full-face adjacentcies is greater than the snake length less one, the new snake is added to the list of valid snakes at process block 1016 and the process loops back to decision block 1011. That is, if the snake passes all of the above tests, it is seen as valid and is added to the list of all valid snakes and a center voxel having a set of “on” neighbors with the given configuration will be seen as possessing sufficient spatial connectivity with that group of neighbors to warrant merging with them. The above-described process continues until the binary number

reaches overflow at decision block 1011, and then the process ends at process block 1020.

[00110] By convention, the neighbors in a 3x3x3 neighborhood have been numbered from 0 to 25, as shown in Fig. 26. Thus, the state of an entire 3x3x3 neighborhood can be captured by a 26-bit number. If the numbers from the array of Fig. 26 represent the bit position in this 26-bit number, then a 1 may be written into each corresponding bit where the voxel is "on," and a 0 may be written into each corresponding bit where the voxel is "off." For example, voxel numbers 0, 1, 2, 3, 4, 6, 7, and 8 are on, and the rest are all off. The 26-bit number describing this neighborhood is shown in the list of Fig. 26 arranged below the array of voxels. Snakes are defined to be the basic unit of spatial connectedness. That is, if all the neighbor voxels of some hypothetical center voxel, defined by a given snake are "on," then the given center voxel is seen as possessing sufficient spatial connectivity to the voxels of the snake to warrant merging with it. That is, there is seen to be sufficient connectivity between the center voxel, and the neighbor voxels within the given snake, that they should be viewed as belonging to the same clump. Snakes are defined by their length (this is configurable, but the present embodiment defines the length of a snake to be 8), by their connectivity (that is, all elements of a snake must be connected to the snake by at least one full-face connection), and by their compactness (snakes are not allowed to be linear, but must wind back on themselves at least once – the degree of the required compactness is configurable).

[00111] It is also possible to use the list of valid snakes to build what is called a "rapid neighborhood assessor." That is, given a hypothetical center voxel, and some neighborhood configuration, it is possible to construct a data structure to allow the determination of which voxels the hypothetical center voxel should be merged with, in essentially one step, that is, instead of sequentially examining each snake. The process of using the list of snakes to construct a rapid neighborhood assessor is shown in Fig. 27. This process includes exhaustively going through every possible neighborhood configuration, determining which snakes are completely contained by each given neighborhood configuration, and recording for every possible neighborhood configuration which neighbors should be merged. The rapid neighborhood assessor, then, includes a list of neighborhood configurations, each expressed as the decimal equivalent of their 26-bit number representation, accompanied by a list of all neighbors with which a hypothetical center voxel should be merged, given the specified neighborhood configuration.

[00112] Fig. 28 shows graphically the process of applying the rapid neighborhood assessor to a hypothetical neighborhood configuration. Referring to Figs. 27 and 28, the process begins by creating binary numbers representing possible neighborhood configurations at process block 1100. The user may construct a binary number from any given neighborhood configuration and use the binary number as an index into a very large array. Each element of this array indicates which neighbors the center voxel should be merged with, given all the possible snakes which might fit into the "on" neighbors (if any). At process block 1110, an iterative process of analyzing the snakes begins by examining the first snake in the list of valid snakes. At decision block 1112, a check is made to determine if the current snake is completely contained by the current neighborhood configuration. If not, the next snake is selected at decision block 1114 and the process loops back to decision block 1112. Once a snake that is completely contained by the current neighborhood configuration is selected, at process block 1116, all of the elements of the current snake are added to the list of neighbors to merge for this neighborhood configuration and the next snake is selected at decision block 1114. This process continues until, at decision block 1114, no more snakes remain. At this point, at decision block 1118, a binary neighborhood description counter is incremented and a check is made to determine if the binary neighborhood description has reached overflow. If not, the process loops back and reiterates until overflow is reached and the process ends at process block 1120.

[00113] The neighborhood configuration, includes "on" and "off" voxels, for example, the one shown in the voxel array in Fig. 28, is converted to its 26-bit binary representation, such as represented in below the voxel array in Fig. 28. The 26-bit binary number defining the neighborhood configuration is then interpreted as a decimal number, and used as the index into the rapid neighborhood assessor. The given index into the rapid neighborhood assessor supplies the numbered list of neighbors which the given center voxel should be merged with, given the specified neighborhood configuration. The process of constructing the rapid neighborhood assessor is extremely computationally intensive; however this needs only to be done once, for a given definition of what a snake is, and subsequently the rapid neighborhood assessor may be loaded from the disk for all subsequent instantiations of the algorithm. Although the rapid neighborhood assessor structure is quite large in terms of its memory utilization requirements, it produces a profound improvement in the speed of execution of the overall algorithm.

[00114] The method is described in the context of reduction of noise/identification of "significant" regions within change detection brain MRI images. The method's ability to identify such regions within noise fields, particularly regions whose amplitude lies below the level of the noise floor, presents possible applications far beyond this domain. The method should, in fact, find application in a large number of image processing fields. Specifically it should find application in any case where subtle regions of signal, which may be below the noise floor but which exhibit spatial connectedness, exist. These include a potential unlimited number of applications, of course, including a range of areas of change detection, such as medical analysis, remote sensing, astronomy, geology, surveillance, classification, and the like. There are many potential applications for the method within medical imaging, but beyond change detection, for example, the algorithm has been used extensively to aid in automated sample point definition, and in lesion finding as described above.

CLAIMS

1. A system for detecting changes in a series of images medical comprising:

an image acquisition system configured to acquire a series of medical images of a region of interest (ROI) in a subject;

an automated sample point generator configured to review the series of medical images and automatically detect portions of the series of medical images indicative of a material expected to be in the ROI and designate sample points therefrom;

a normalizer configured to utilize the sample points to normalize intensities within the series of medical images;

an image registration system configured to utilize the sample points to automatically register the series of images such that common areas within the ROI of the subject are aligned; and

a change tracker configured to automatically identify changes between the series of medical images and display an indication of the identified changes.

2. The system of claim 1 wherein the automated sample point generator is further configured to create a mathematical description of properties of abnormalities appearing in a particular image in the series of medical images.

3. The system of claim 2 wherein the change tracker is further configured to utilize the mathematical description of the properties of abnormalities appearing in a particular image to further identify changes between the images.

4. The system of claim 2 wherein the abnormalities include lesions and enhancing lesions.

5. The system of claim 1 further comprising a mask generator configured to generate a parenchyma mask from the registered images and wherein the change tracker is further configured to utilize the parenchyma mask to identify changes between the series of medical images.

6. The system of claim 1 wherein the normalizer utilizes a mathematical model of partial character/partial volume effects to normalize intensities within the series of medical images.

7. The system of claim 1 further comprising a significant region detector configured to automatically separates the regions of actual signal from background noise in the series of medical images.

8. The system of claim 1 wherein the series of medical images include at least one of MRI images and CT images.

9. The system of claim 1 wherein the material indicative of a material expected to be in the ROI includes cerebral spinal fluid (CSF), white matter, gray matter, bone, and lesions.

10. A method for detecting changes in a series of images medical comprising the steps of:

- a) acquiring a series of medical images of a region of interest (ROI) in a subject;
- b) automatically reviewing the series of medical images to detect portions of the series of medical images indicative of a material expected to be in the ROI and designate sample points therefrom;
- c) normalizing intensities within the series of medical images using the sample points;
- d) automatically registering the series of images such that common areas within the ROI of the subject are aligned using the sample points as references; and
- e) indicating changes between the series of medical images.

11. The method of claim 10 wherein step b) includes creating a mathematical description of properties of abnormalities appearing in a particular image in the series of medical images.

12. The method of claim 10 wherein step b) further includes utilizing the mathematical description of the properties of lesions and enhancing lesions appearing in a particular image to further identify changes between the images.

13. The method of claim 10 wherein step d) further includes generating a parenchyma mask from the registered images and wherein step e) includes utilizing the parenchyma mask to identify changes between the series of medical images.

14. The method of claim 10 wherein step c) includes utilizing a mathematical model of partial character/partial volume effects.

15. A system for detecting changes in a series of images medical comprising:

- an automated sample point generator configured to review a series of medical images and automatically detect portions of the series of medical images indicative of a material expected to be in a ROI included in each of the series of medical images and designate sample points within the series of medical images therefrom;

- a normalizer configured to utilize the sample points to normalize intensities within the series of medical images;

- a lesion finder configured to identify potential lesions and enhancing lesions within the series of medical images;

- an image registration system configured to utilize the sample points to automatically register the series of images such that common areas within the ROI of the subject are aligned; and

- a change tracker configured to automatically identify changes between the series of medical images and display an indication of the identified changes.

16. The system of claim 15 further comprising a noise reducer configured to identify spatially cohesive regions in the series of medical images embedded in a zero-mean field of noise.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

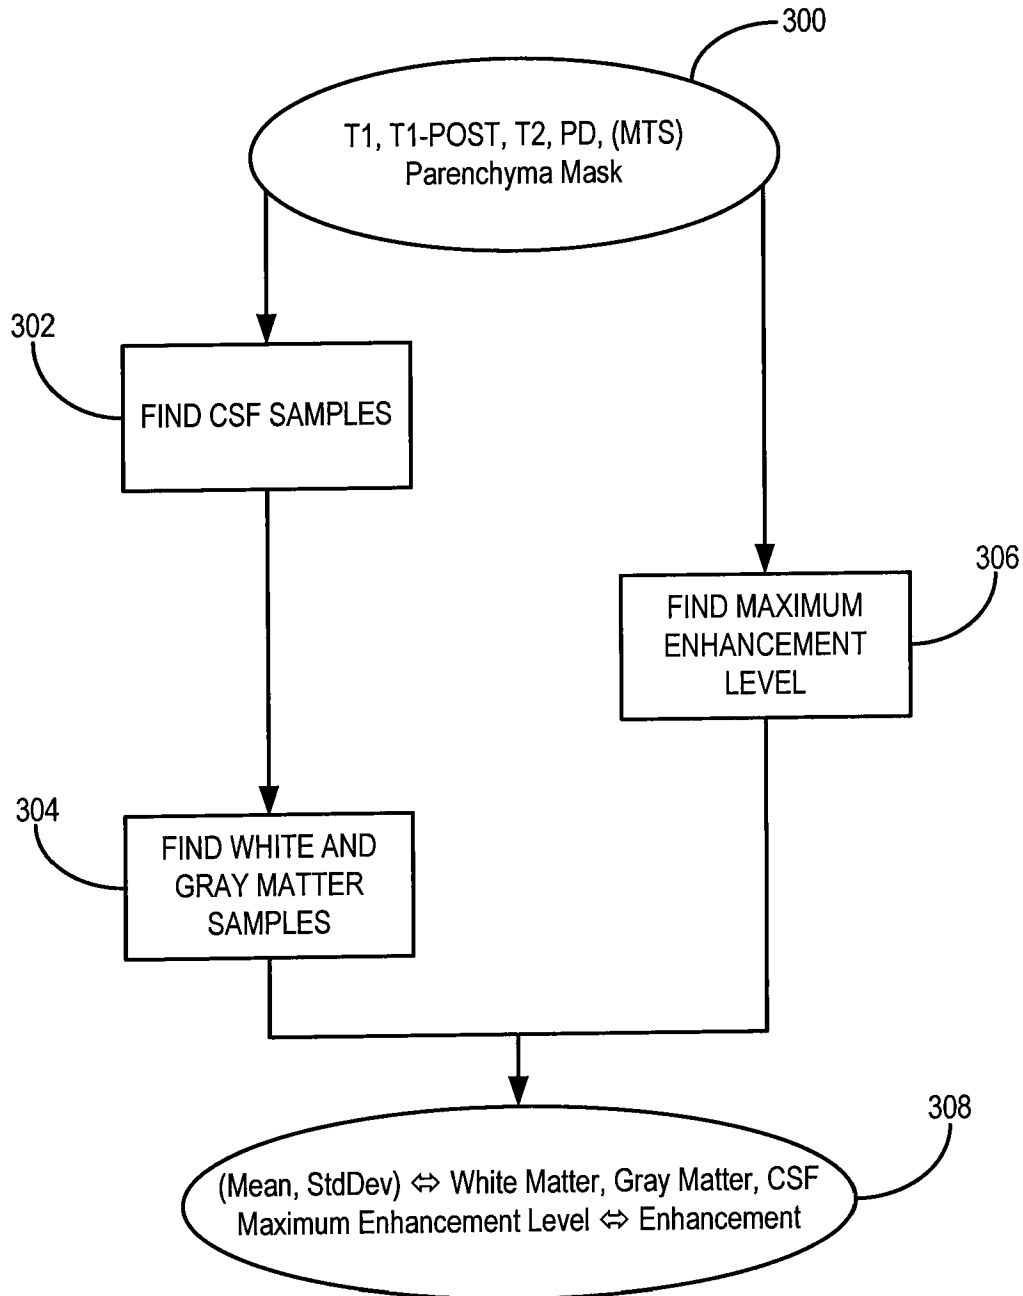


FIG. 5

5/21

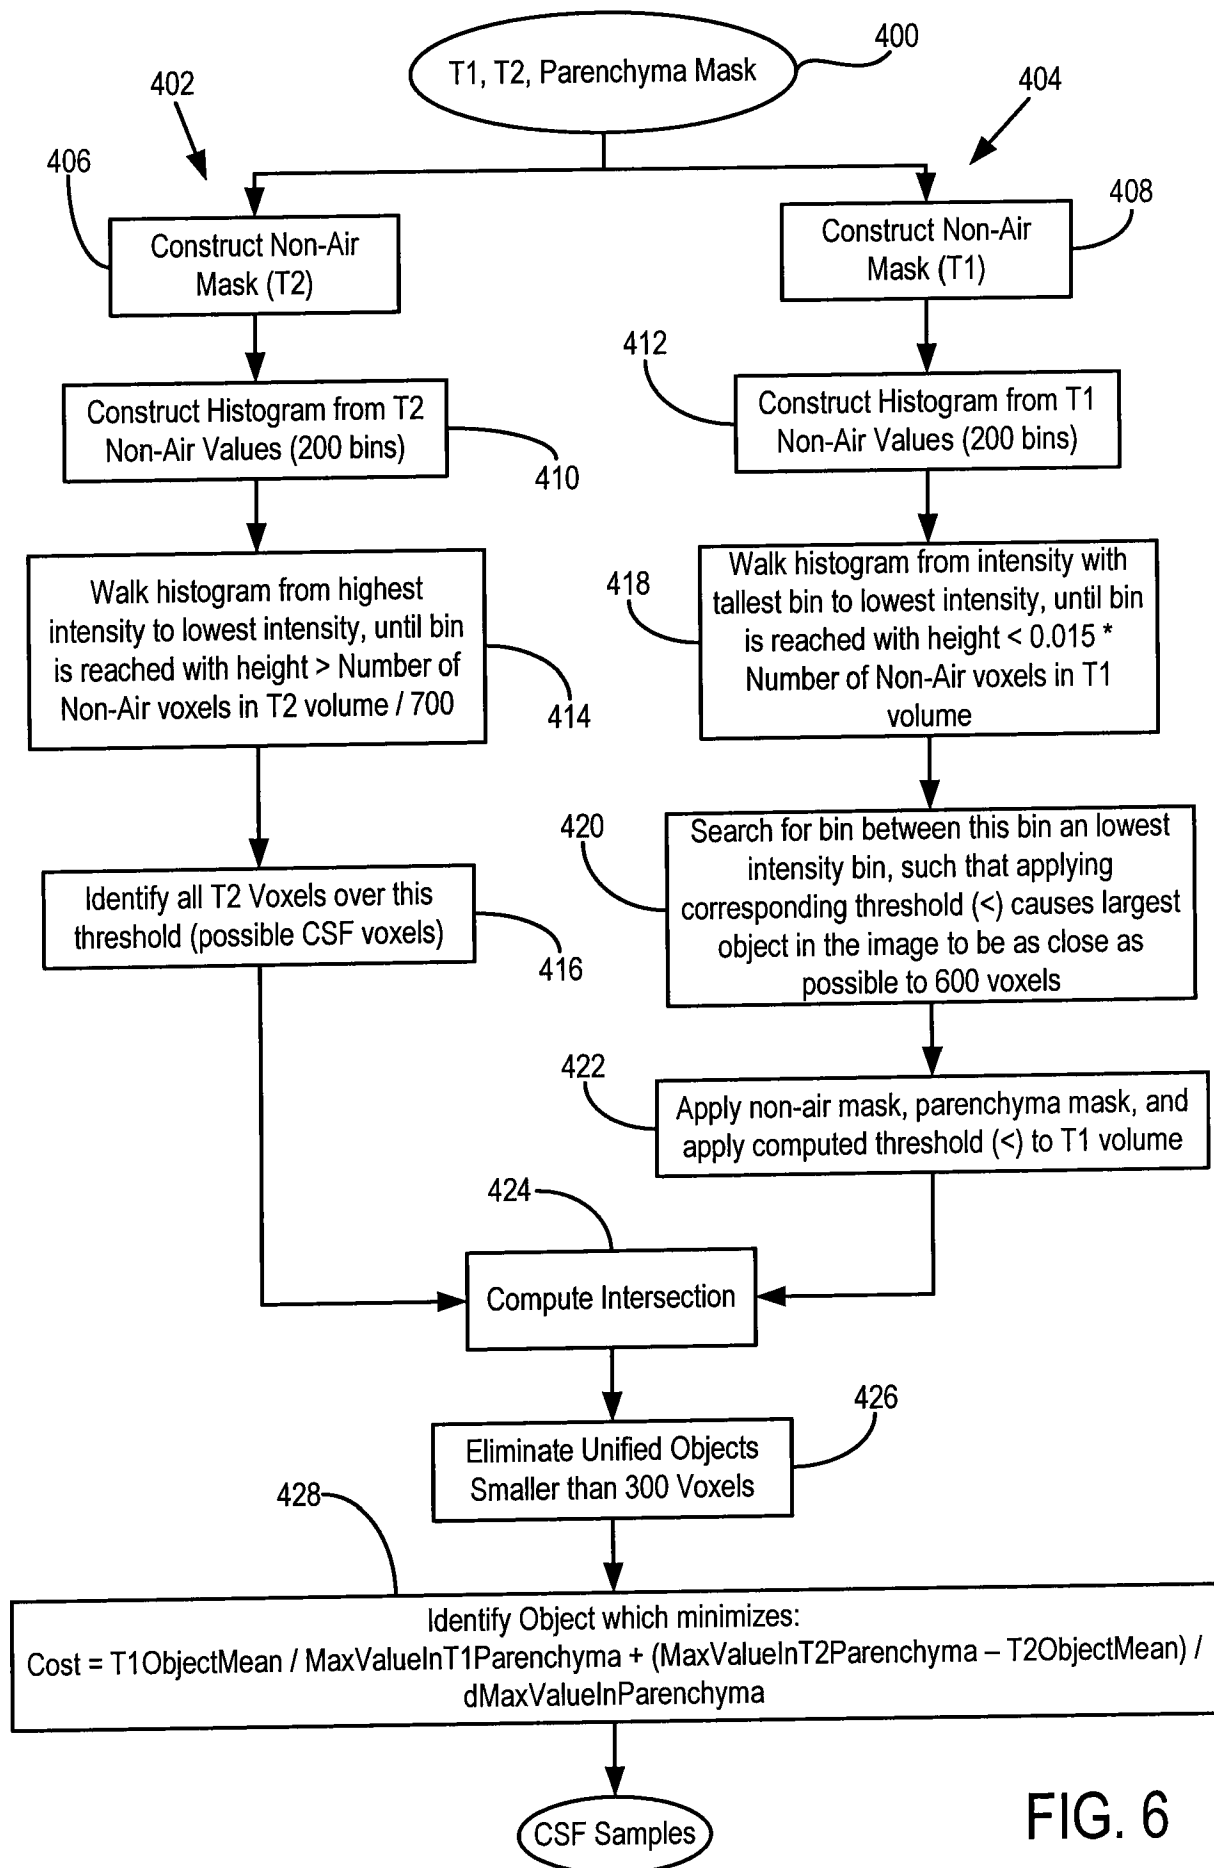


FIG. 6

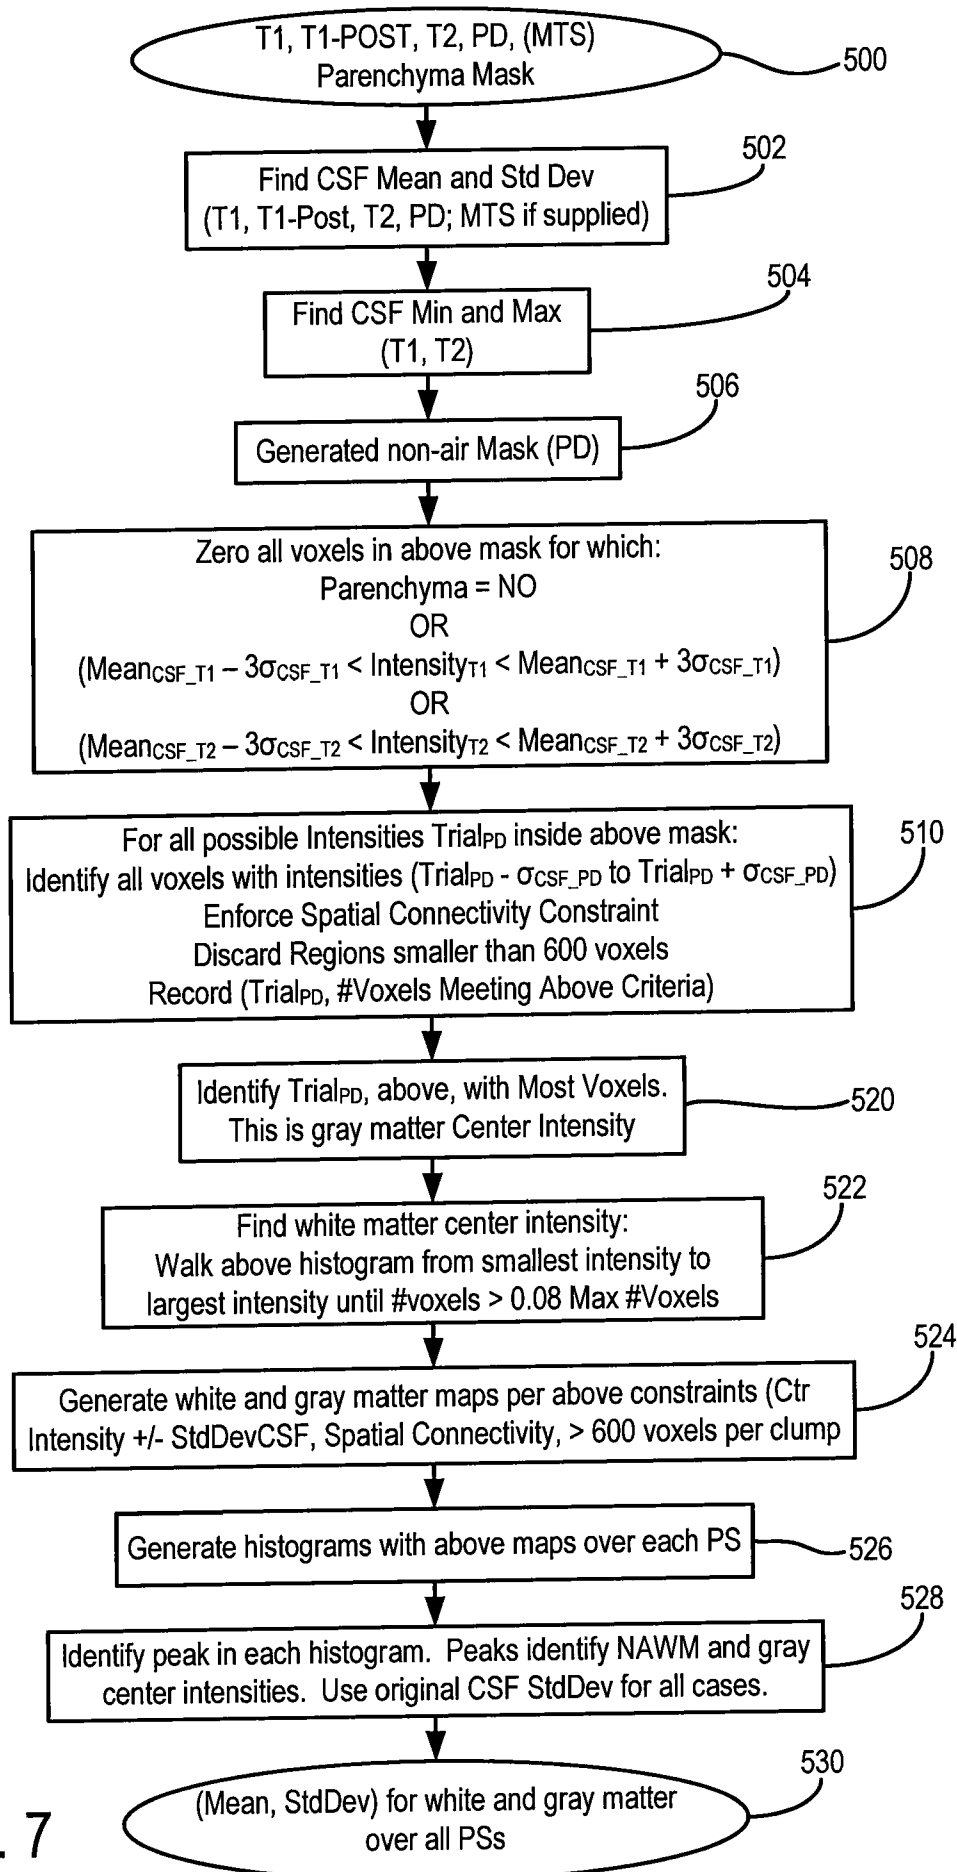


FIG. 7




FIG. 8

8/21

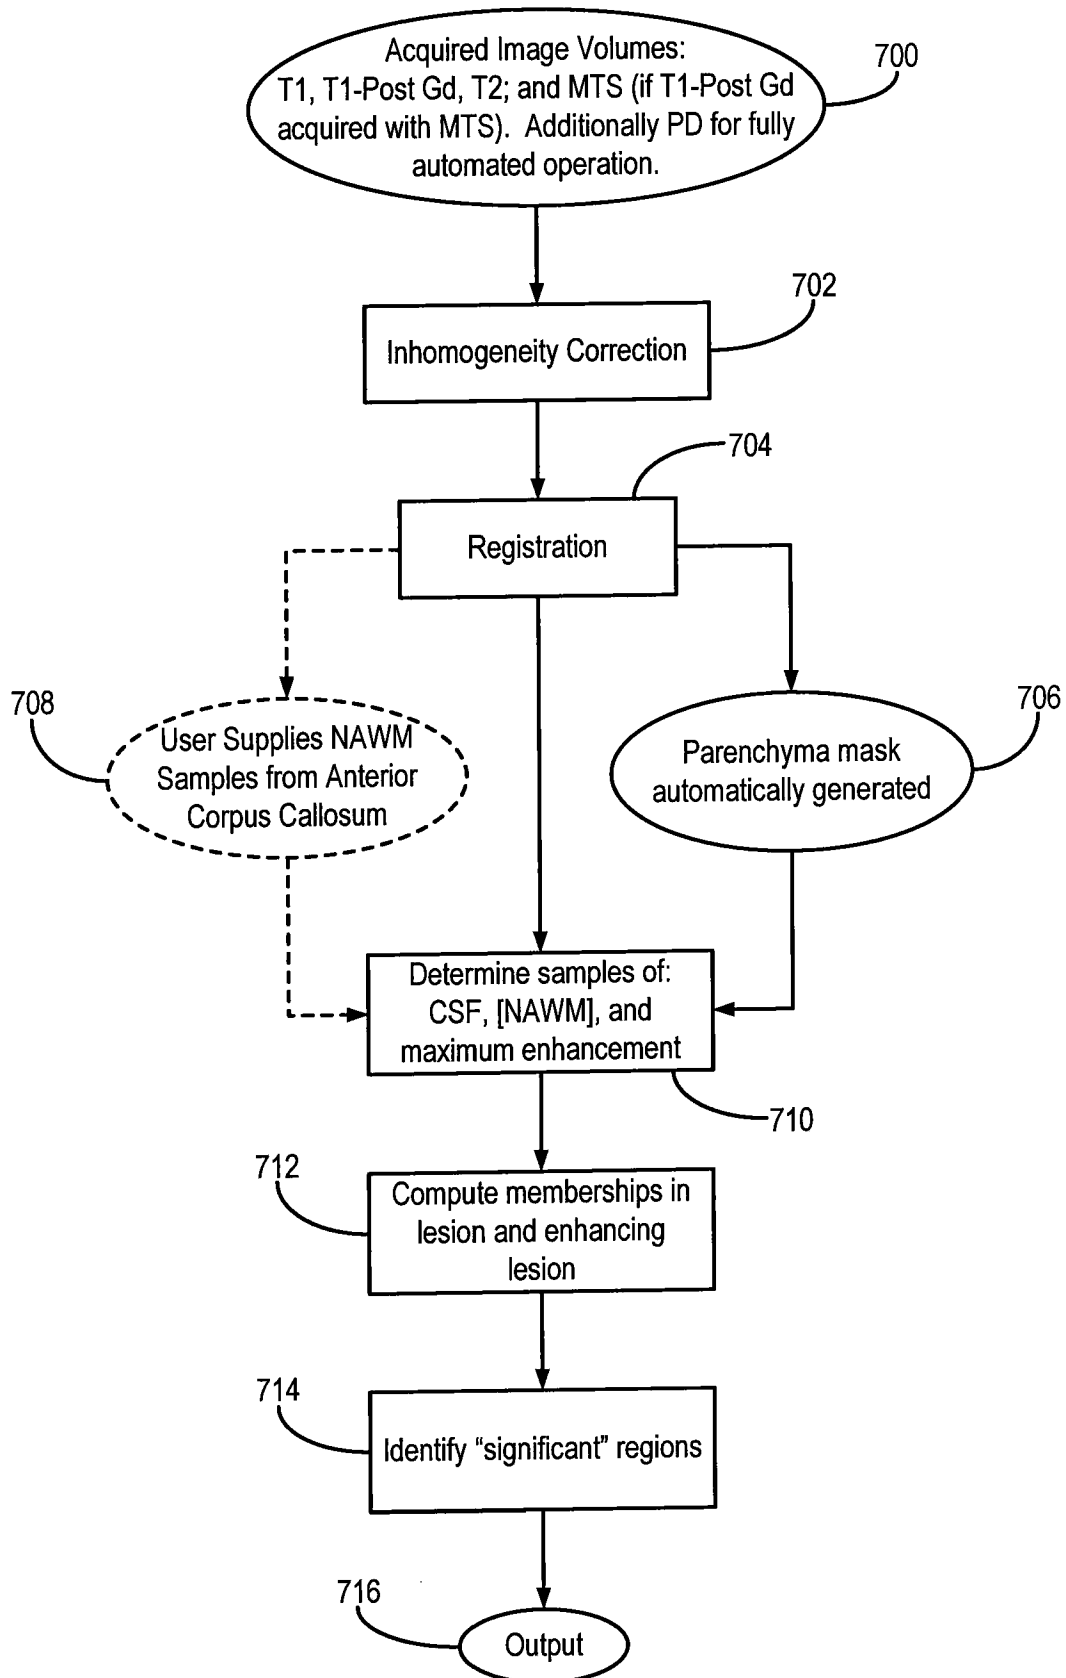


FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15

12/21

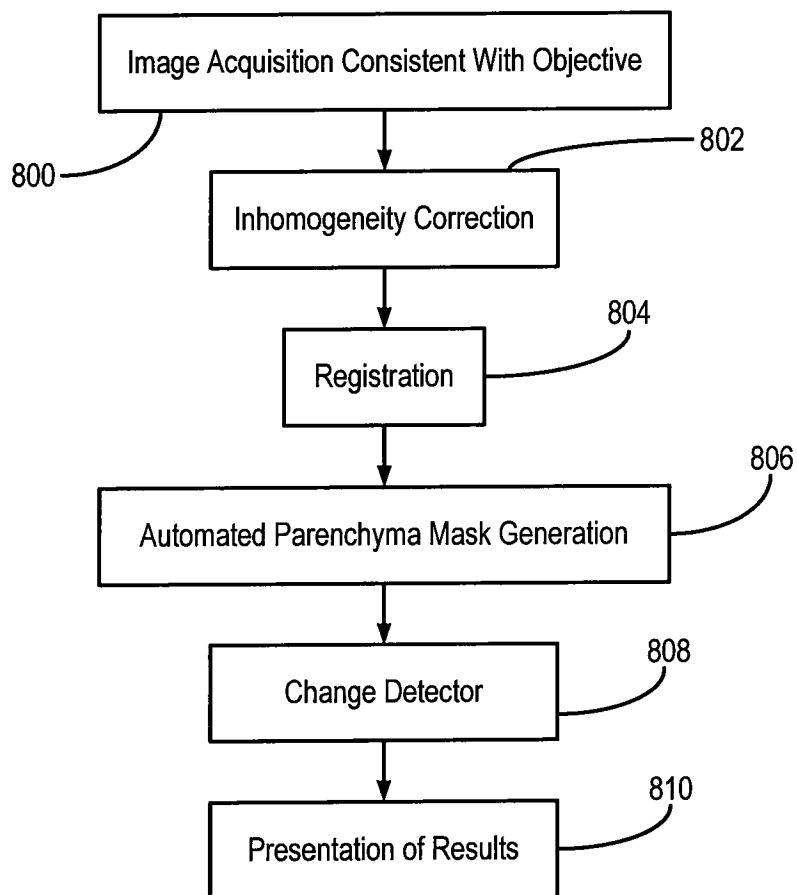


FIG. 16

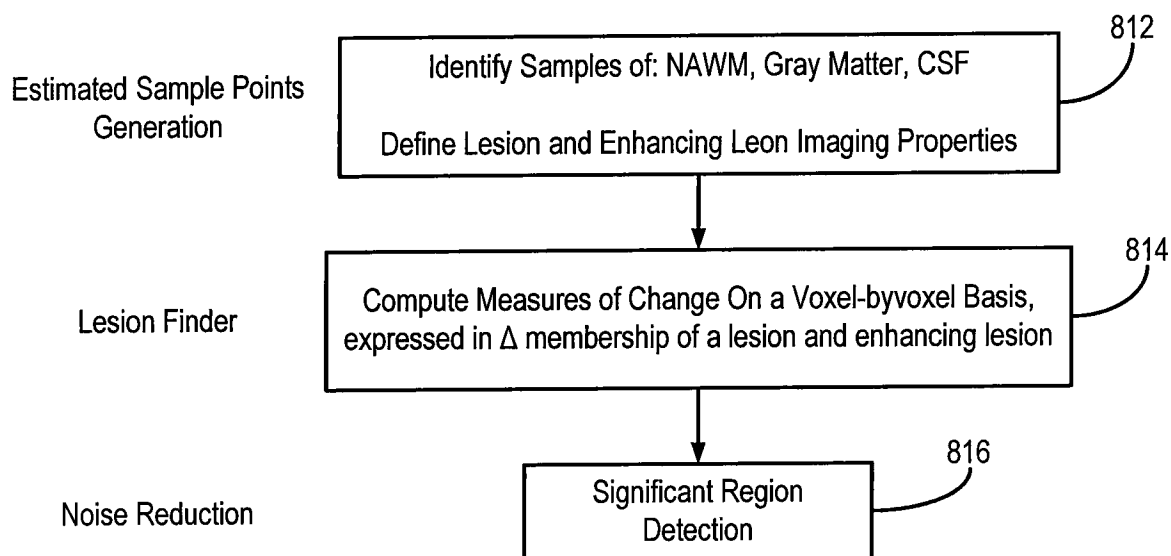


FIG. 17




FIG. 18




FIG. 19




FIG. 20




FIG. 21

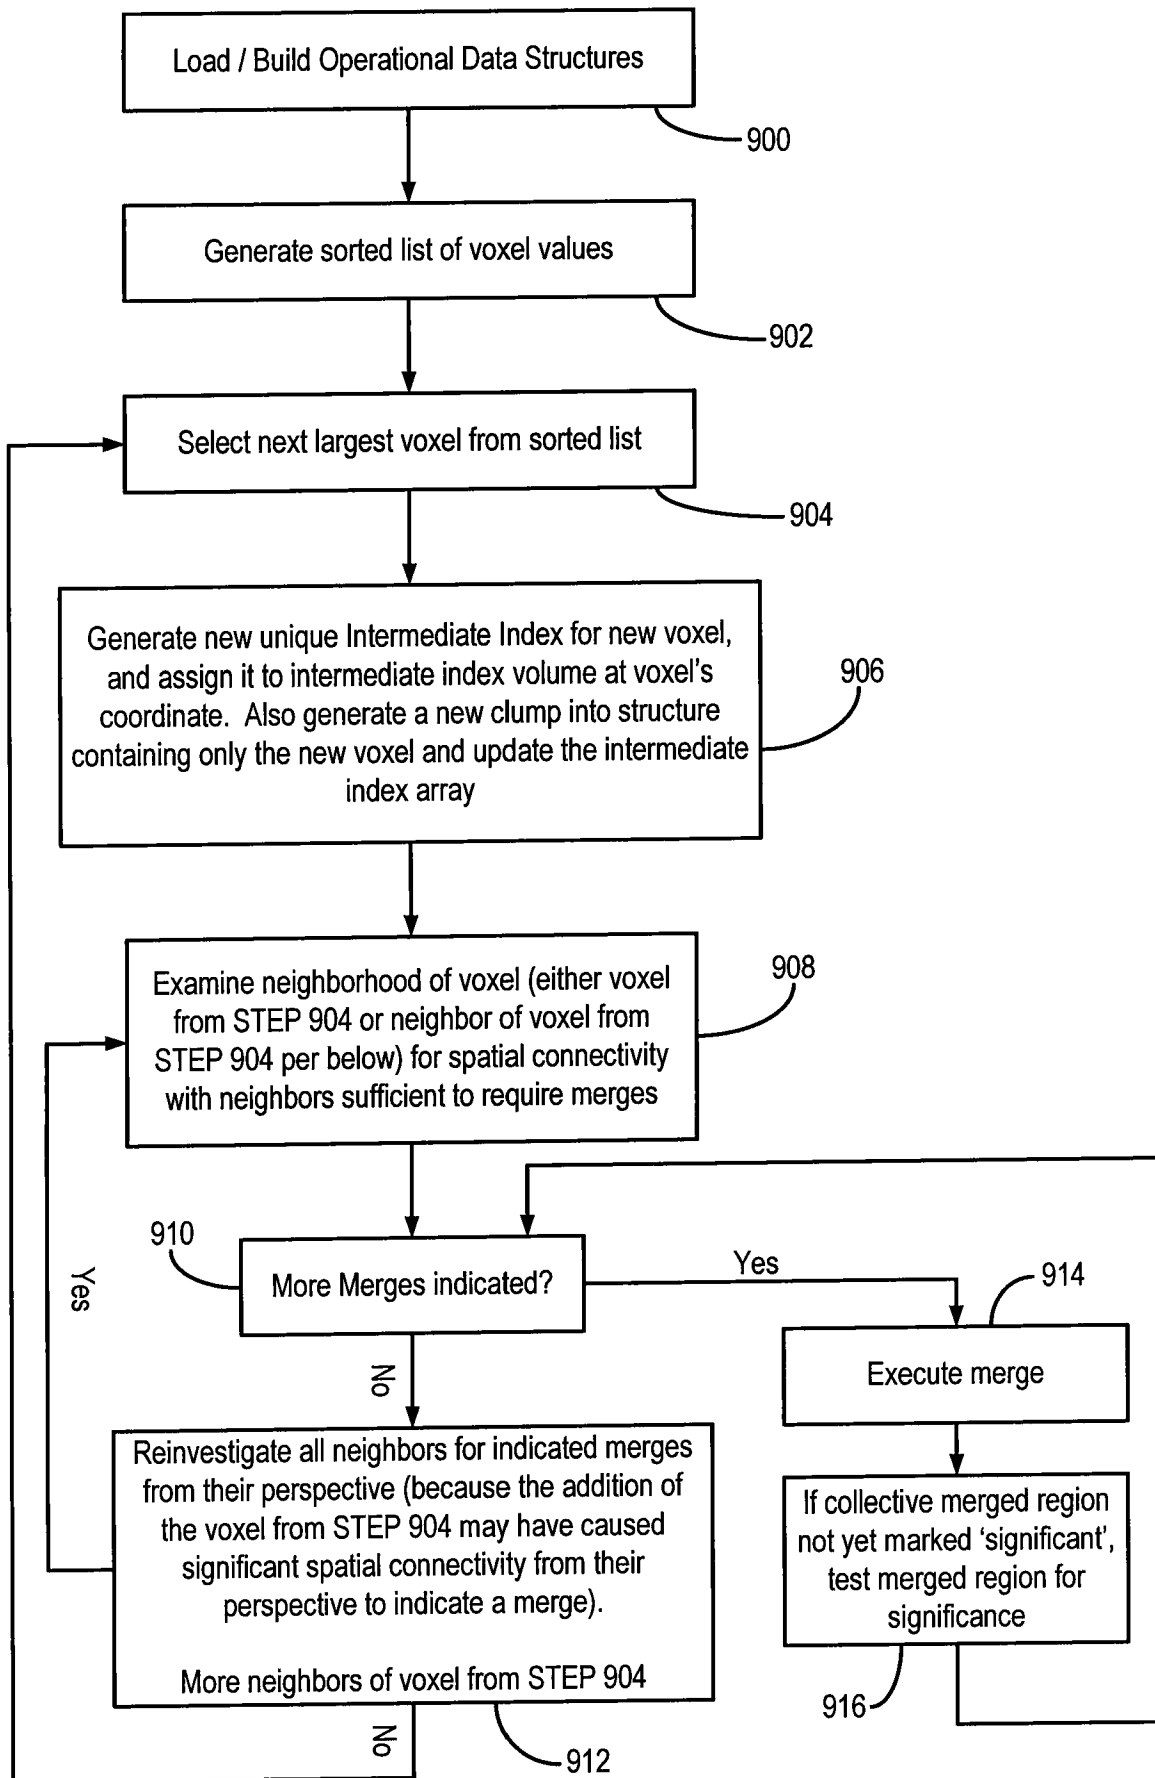


FIG. 22

Clump mean required for significance (expressed in
number of std. devs. of noise)


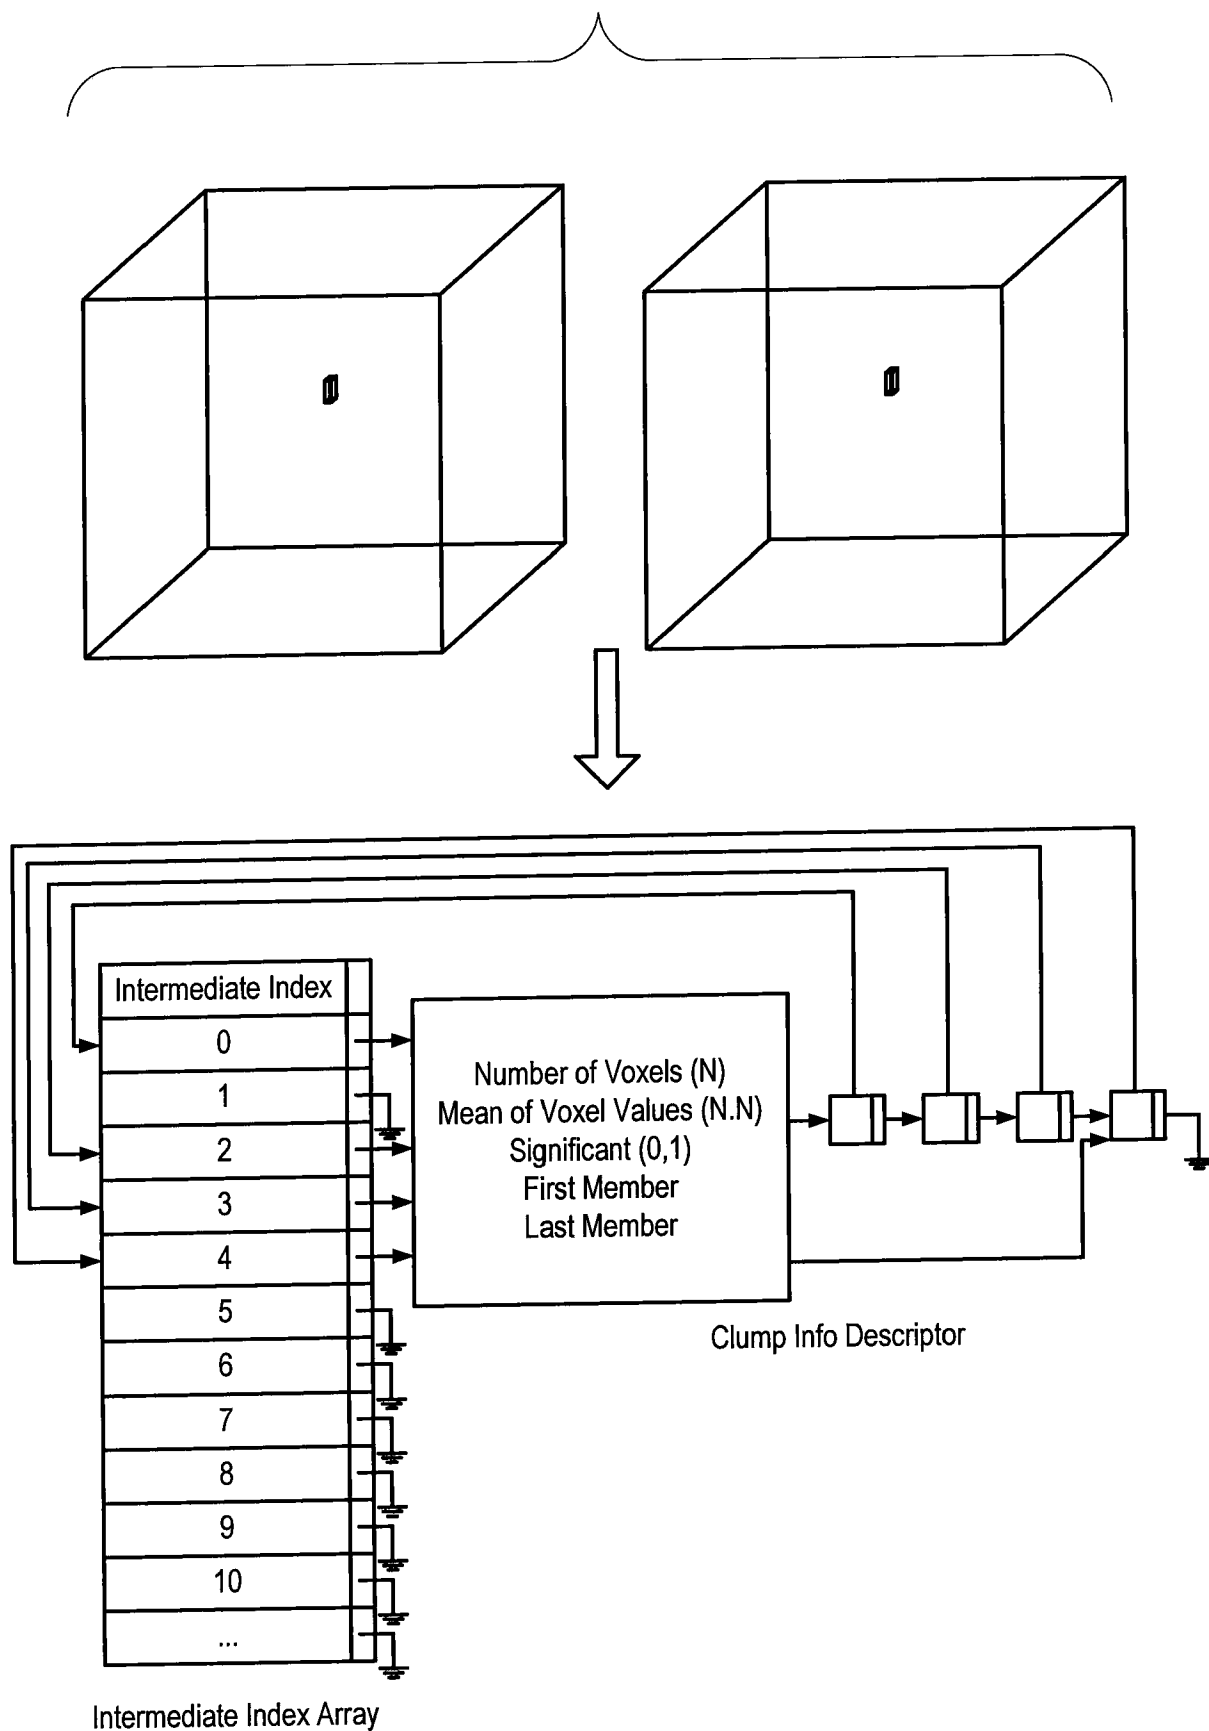


FIG. 23

FIG. 24







FIG. 25



0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	0	1	1	1	1	1
25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

FIG. 26

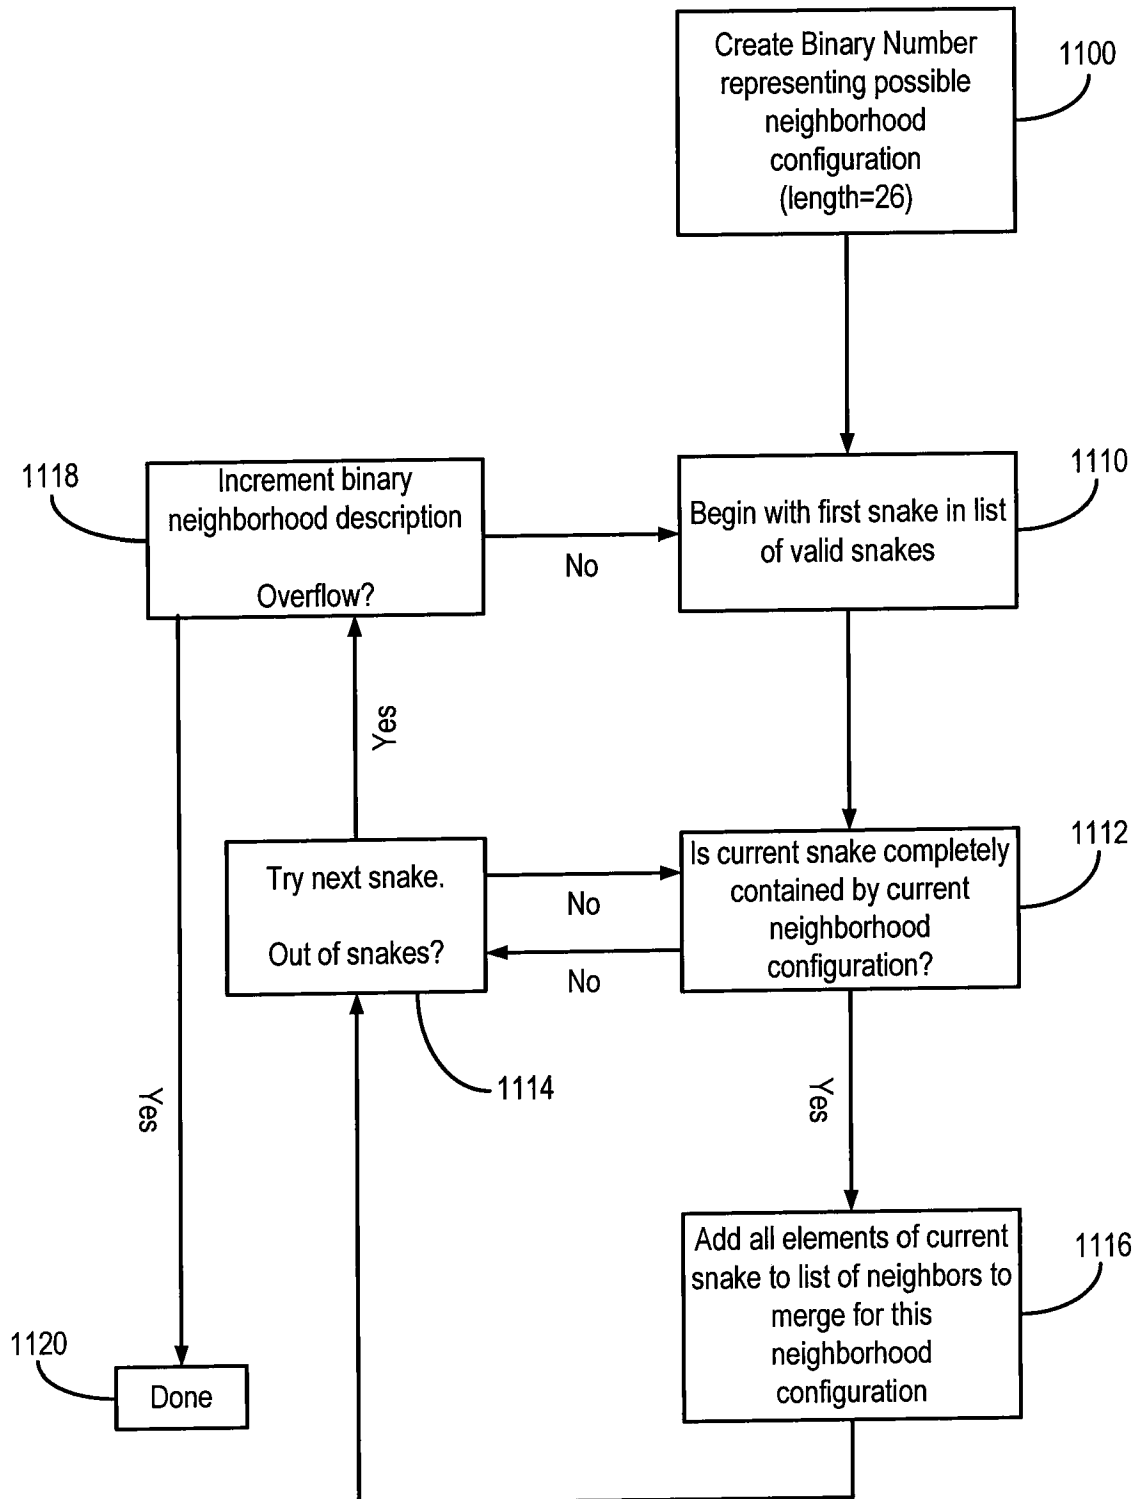


FIG. 27




FIG. 28

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 08/68826

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06K 9/00 (2008.04)

USPC - 382/128

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06K 9/00 (2008.04); 382/128Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
382/130

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Google; Google Scholar; Google Patents; PubWest(PGPB,USPT,USOC,EPAB,JPAB);

Search Terms Used: nuclear magnetic resonance imaging, MRI, NMR, track changes, detect changes, images, series, medical, sample, normalize, indicate change, identify change, abnormal, parenchyma mask, region

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2006/0127880 A1 (HARRIS et al.) 15 June 2006 (15.06.2006), para [0005]-[0018], para [0031]-[0080], para [0319]	1-16
Y	US 2006/0088836 A1 (WOHLGEMUTH et al.) 27 April 2006 (27.04.2006), para [0263], para [0338]-[0393], para [0495], para [0620]	1-16
A	US 6,272,370 B1 (GILLIES et al.) 07 August 2001 (07.08.2001)	1-16
A	US 7,048,716 B1 (KUCHARCZYK et al.) 23 May 2006 (23.05.2006)	1-16

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

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

13 October 2008 (13.10.2008)

Date of mailing of the international search report

20 OCT 2008

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774