



(19) **United States**

(12) **Patent Application Publication**
Beckers et al.

(10) **Pub. No.: US 2018/0256042 A1**
(43) **Pub. Date: Sep. 13, 2018**

(54) **MEDICAL IMAGING AND EFFICIENT SHARING OF MEDICAL IMAGING INFORMATION**

Publication Classification

(71) Applicant: **ARTERYS INC.**, San Francisco, CA (US)

(51) **Int. Cl.**
A61B 5/026 (2006.01)
A61B 5/00 (2006.01)
G16H 30/40 (2006.01)
G06T 7/00 (2006.01)
G06T 7/269 (2006.01)

(72) Inventors: **Fabien Beckers**, San Francisco, CA (US); **John Axerio-Cilies**, San Francisco, CA (US); **Torin Arni Taerum**, Calgary (CA); **Albert Hsiao**, Menlo Park, CA (US); **Giovanni DE FRANCESCO**, Calgary (CA); **Darryl BIDULOCK**, Calgary (CA); **Tristan Jugdev**, Calgary (CA); **Robert Newton**, Calgary (CA)

(52) **U.S. Cl.**
CPC *A61B 5/0263* (2013.01); *A61B 5/0013* (2013.01); *A61B 5/0044* (2013.01); *A61B 5/7203* (2013.01); *G06T 2207/30104* (2013.01); *G06T 7/0012* (2013.01); *G06T 7/269* (2017.01); *A61B 2576/023* (2013.01); *G06T 2207/10088* (2013.01); *G16H 30/40* (2018.01)

(21) Appl. No.: **15/779,447**

(22) PCT Filed: **Nov. 29, 2016**

(86) PCT No.: **PCT/US2016/064031**

§ 371 (c)(1),

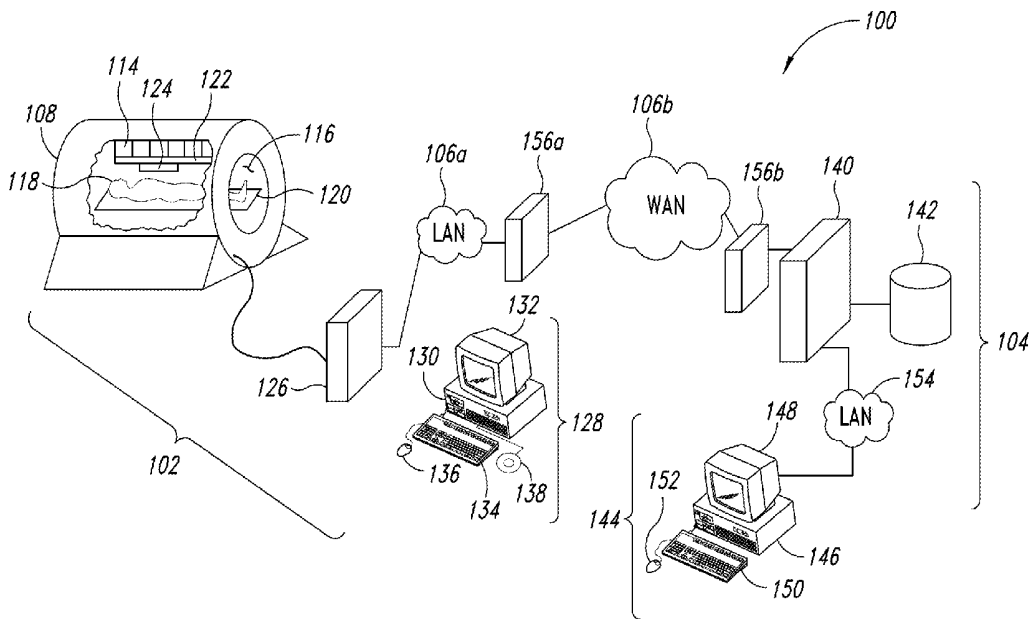
(2) Date: **May 25, 2018**

ABSTRACT

(57) An MRI image processing and analysis system may identify instances of structure in MRI flow data, e.g., coherency, derive contours and/or clinical markers based on the identified structures. The system may be remotely located from one or more MRI acquisition systems, and perform: perform error detection and/or correction on MRI data sets (e.g., phase error correction, phase aliasing, signal unwrapping, and/or on other artifacts); segmentation; visualization of flow (e.g., velocity, arterial versus venous flow, shunts) superimposed on anatomical structure, quantification; verification; and/or generation of patient specific 4-D flow protocols. An asynchronous command and imaging pipeline allows remote image processing and analysis in a timely and secure manner even with complicated or large medical imaging data sets and metadata.

Related U.S. Application Data

(60) Provisional application No. 62/415,203, filed on Oct. 31, 2016, provisional application No. 62/260,565, filed on Nov. 29, 2015.



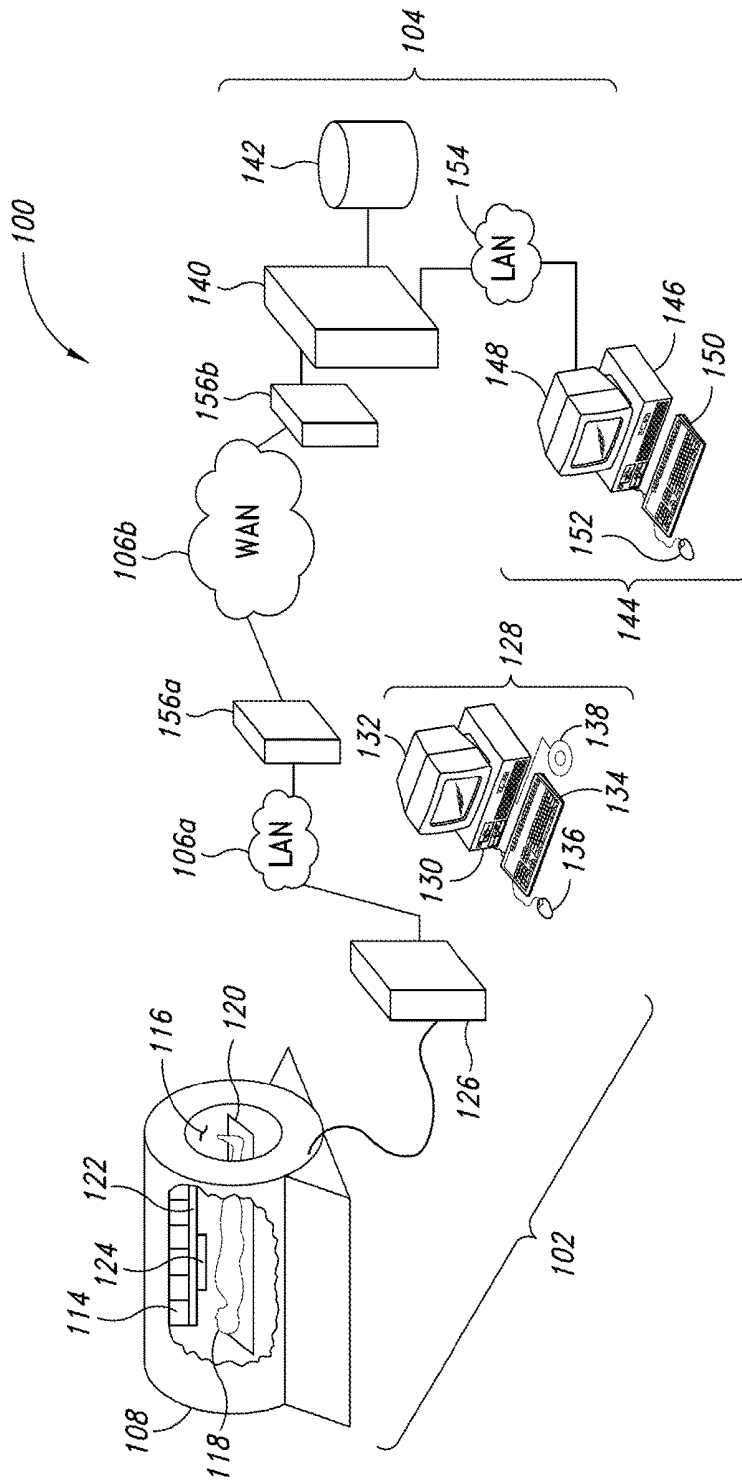


FIG. 1

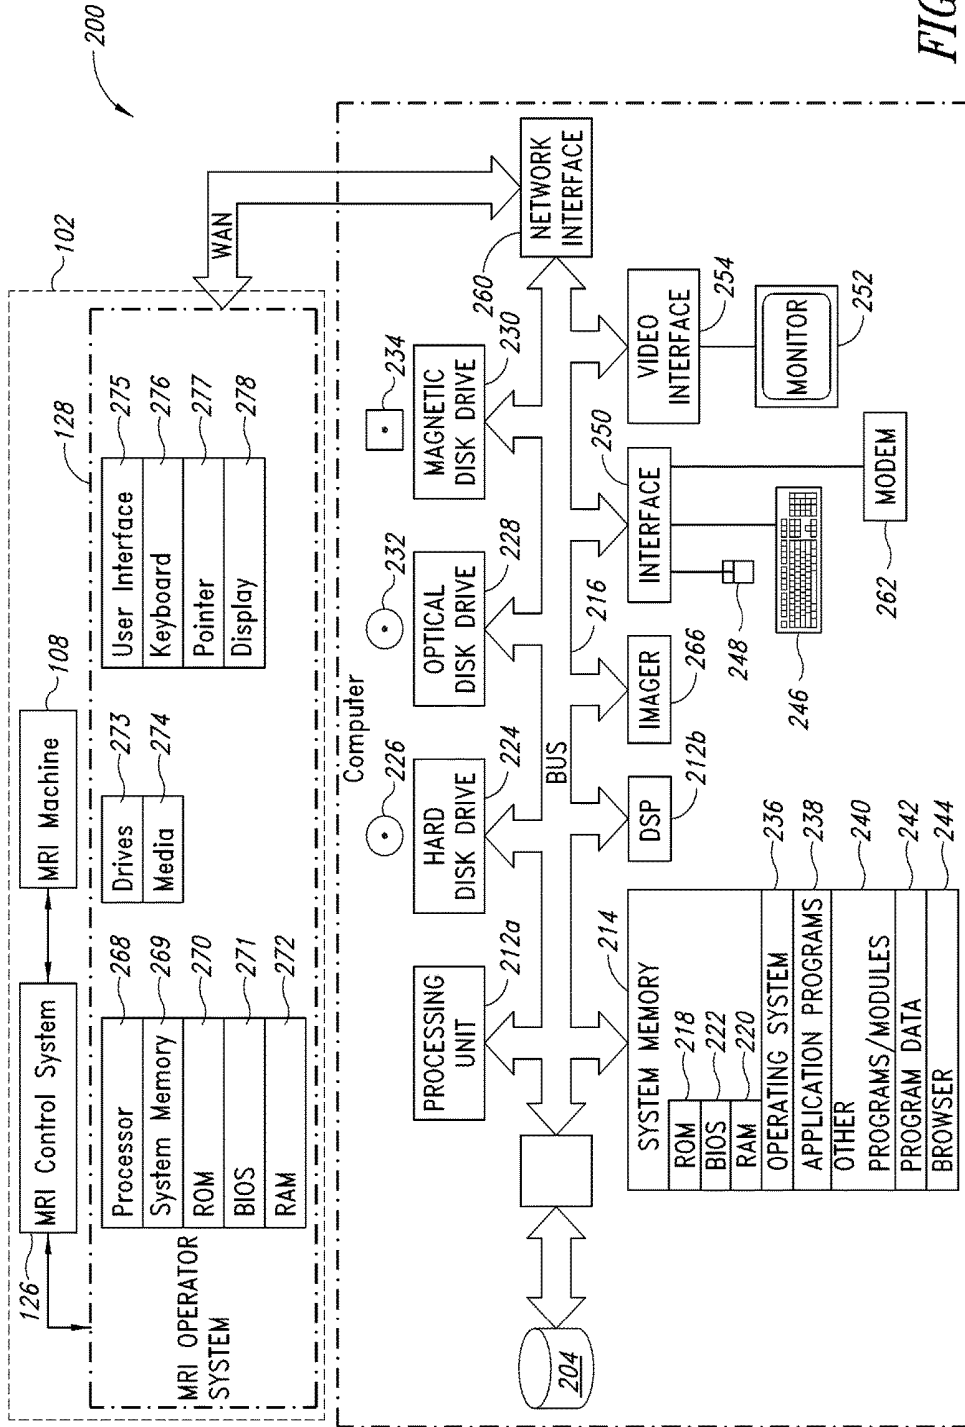


FIG. 2

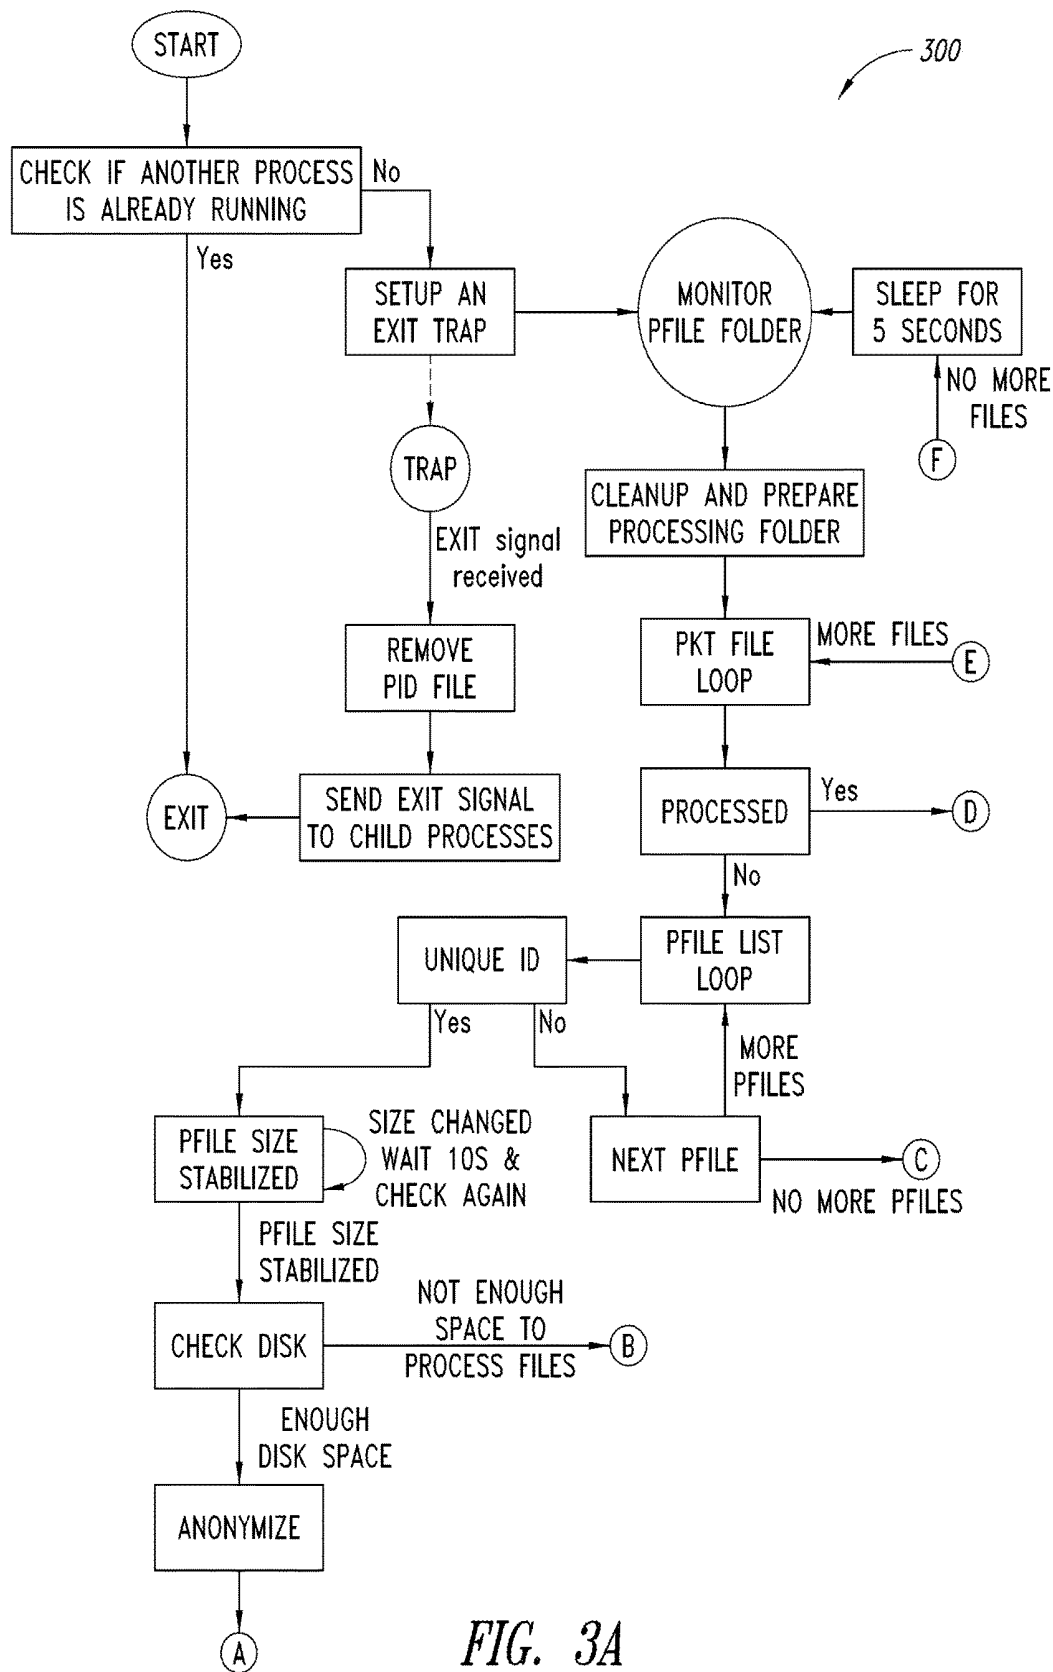


FIG. 3A

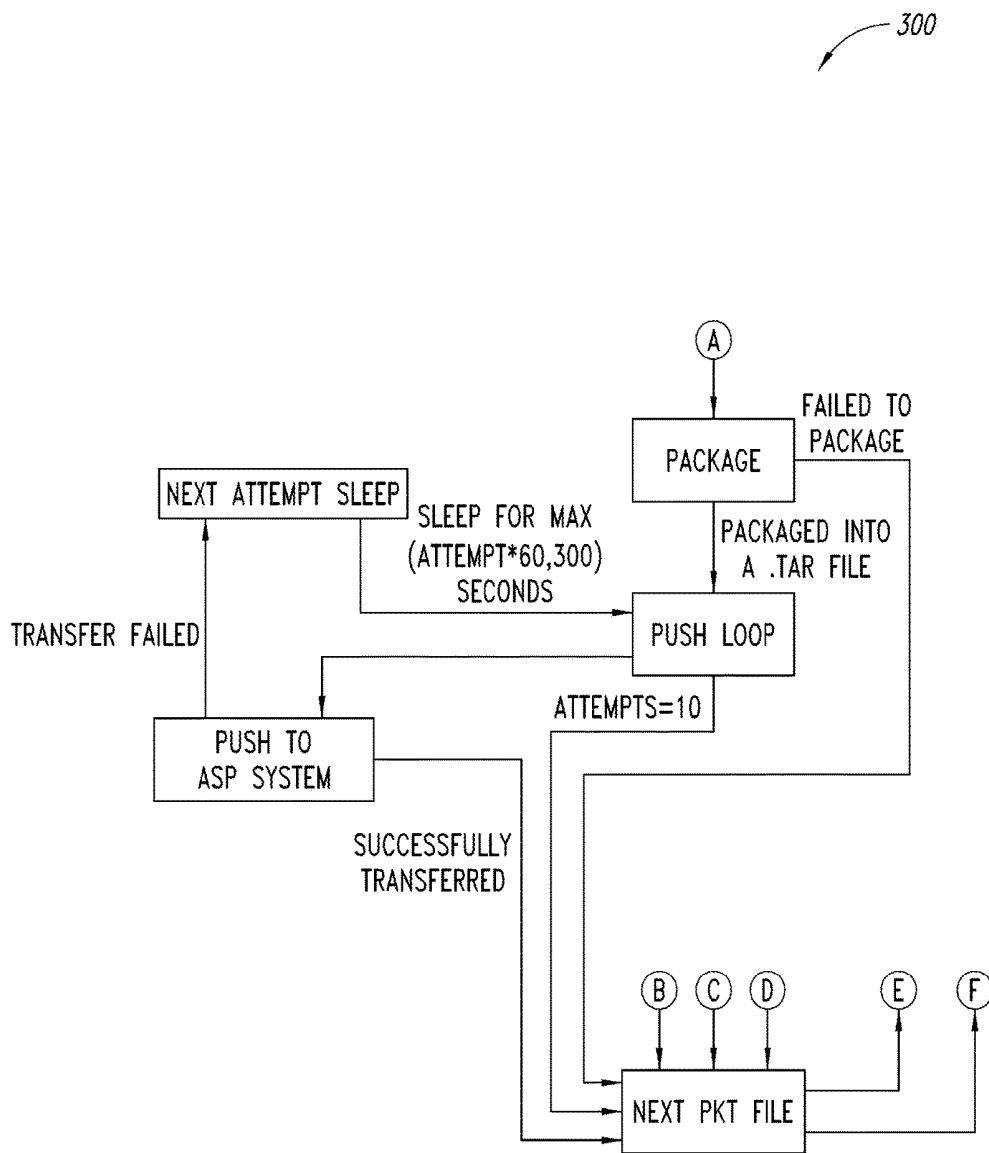


FIG. 3B

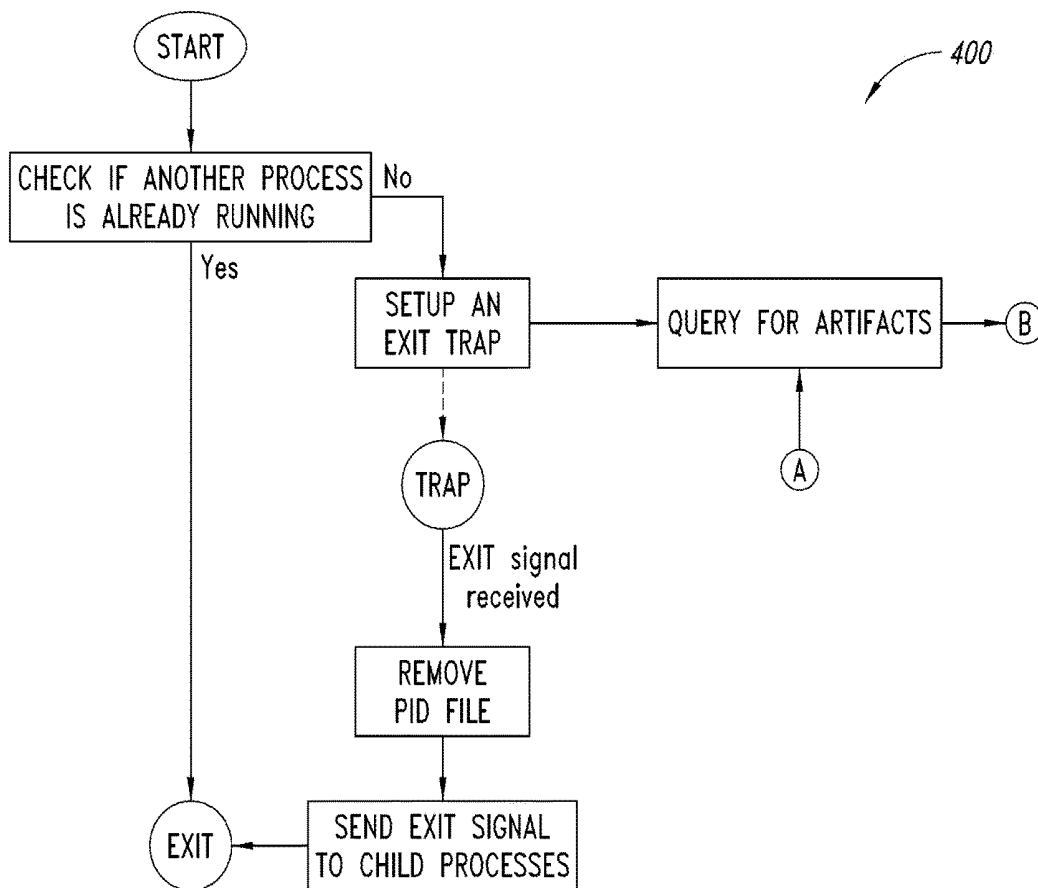


FIG. 4A

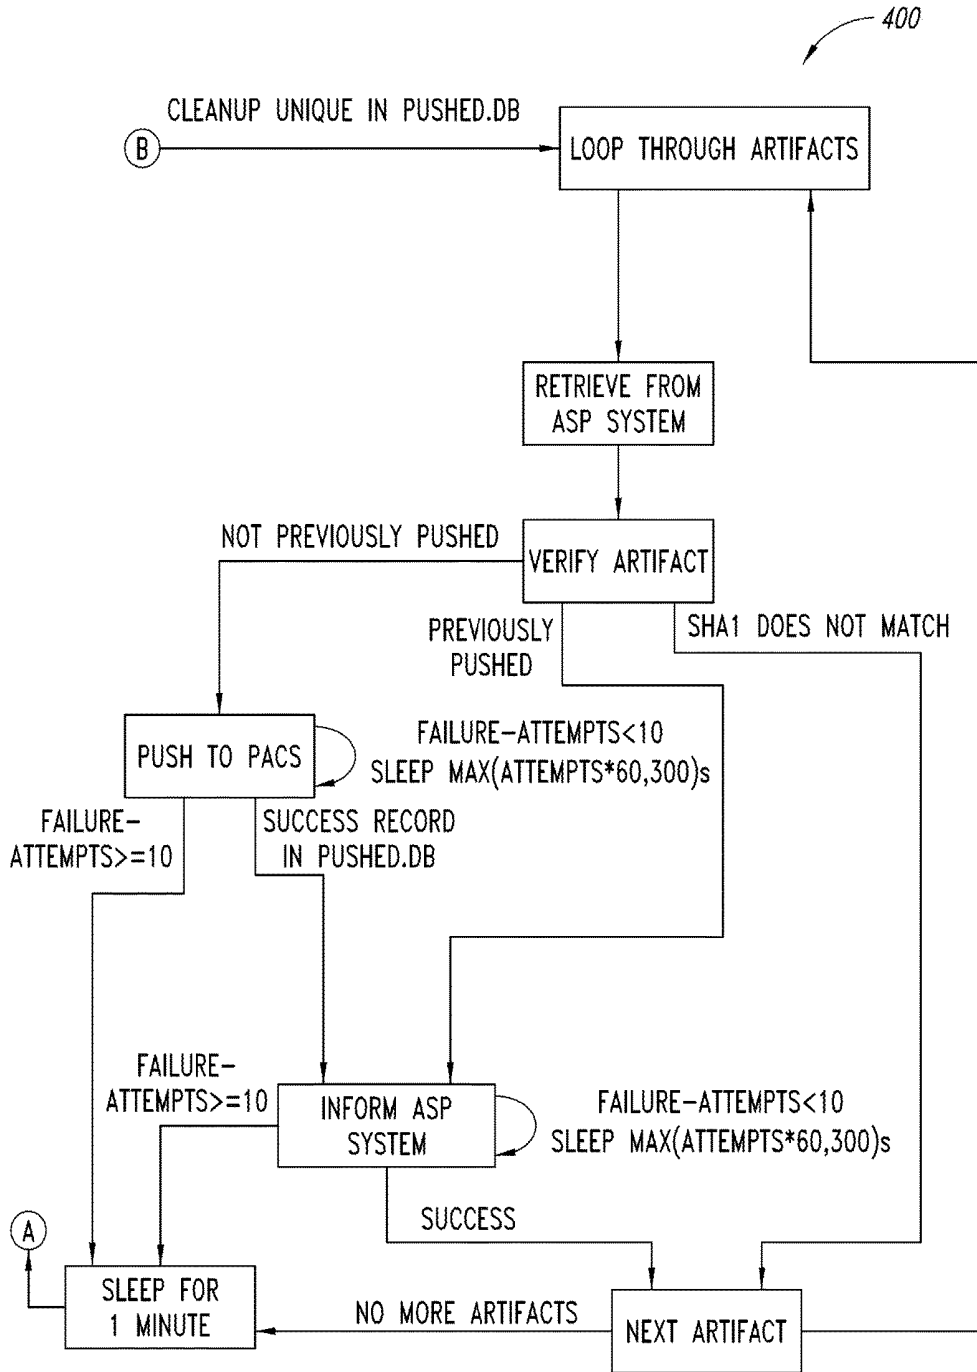


FIG. 4B

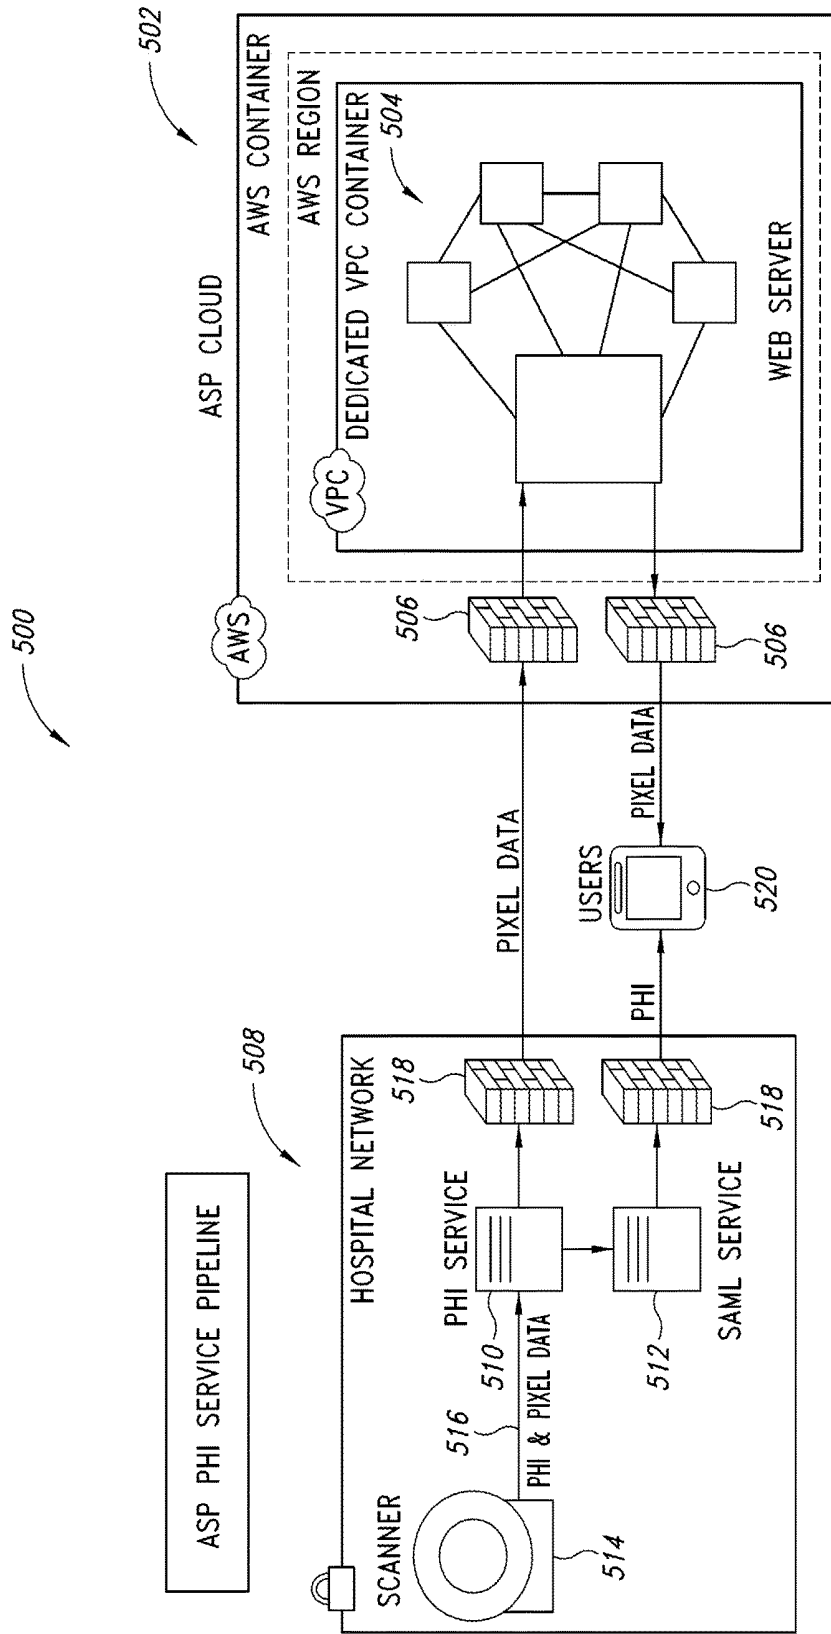


FIG. 5

MEDICAL IMAGING AND EFFICIENT SHARING OF MEDICAL IMAGING INFORMATION

BACKGROUND

Technical Field

[0001] The present disclosure generally relates to magnetic resonance imaging (MRI), for instance four-dimensional (4D) flow MRI, and the sharing of medical imaging and other information over communications networks or channels.

Description of the Related Art

[0002] MRI is most commonly employed in medical imaging, although can be used in other fields. MRI machines include a main magnet which is typically an annular array of coils having a central or longitudinal bore. The main magnet is capable of producing a strong stable magnetic field (e.g., 0.5 Tesla to 3.0 Tesla). The bore is sized to receive at least a portion of an object to be imaged, for instance a human body. When used in medical imaging applications, the MRI machine may include a patient table which allows a prone patient to be easily slid or rolled into and out of the bore.

[0003] MRI machines also include gradient magnets. The gradient magnets produce a variable magnetic field that is relatively smaller than that produced by the main magnet (e.g., 180 Gauss to 270 Gauss), allowing selected portions of an object (e.g., patient) to be imaged. MRI machines also include radio frequency (RF) coils which are operated to apply radiofrequency energy to selected portions of the object (e.g., patient) to be imaged. Different RF coils may be used for imaging different structures (e.g., anatomic structures). For example, one set of RF coils may be appropriate for imaging a neck of a patient, while another set of RF coils may be appropriate for imaging a chest or heart of the patient. MRI machines commonly include additional magnets, for example resistive magnets and/or permanent magnets.

[0004] The MRI machine typically includes, or is communicatively coupled to a computer system used to control the magnets and/or coils and/or to perform image processing to produce images of the portions of the object being imaged. Conventionally, MRI machines produce magnitude data sets which represent physical structures, for instance anatomical structures. The data sets are often conform to the Digital Imaging and Communications in Medicine (DICOM) standard. DICOM files typically include pixel data and metadata in a prescribed format.

BRIEF SUMMARY

[0005] Various approaches that improve 4D flow imaging are described herein.

[0006] A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically correct phase aliasing may be summarized as including receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data including respective anatomical structure and velocity for each of a plurality of voxels; for each of at least some of the plurality of voxels, identifying, by the at least one processor, sharp gradients in flow velocity near a velocity encoding parameter; connecting, by the at least one

processor, all of the voxels identified as having a sharp gradient to define an enclosed boundary; determining, by the at least one processor, whether all of the voxels in the enclosed boundary are aliased; and responsive to determining that all of the voxels in the enclosed boundary are aliased: adding, by the at least one processor, a multiple of the velocity encoding parameter to the velocity for each of the voxels in the enclosed boundary; or subtracting, by the at least one processor, a multiple of the velocity encoding parameter to the velocity for each of the voxels in the enclosed boundary.

[0007] The method may further include responsive to determining that not all of the voxels in the enclosed boundary are aliased, analyzing, by the at least one processor, the respective velocities of neighboring voxels in the enclosed boundary; and modifying, by the at least one processor, the velocity of each of the voxels in the enclosed boundary based at least in part on the analyzed respective velocities of the neighboring voxels. Modifying the velocity of each of the voxels in the enclosed boundary may include modifying the velocity of each of the voxels by an amount determined to minimize the discontinuity across the neighboring voxels. Determining whether all of the voxels in the enclosed boundary are aliased may include determining whether there are any sharp gradients between neighboring voxels in the enclosed boundary.

[0008] A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system may be summarized as including receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data including respective anatomical structure and velocity information for each of a plurality of voxels; defining, by the at least one processor, a volume within the anatomical structure; minimizing, by the at least one processor, a divergence of a velocity field for at least some of the plurality of voxels that represent the volume within the anatomical structure; and correcting, by the at least one processor, velocity information for the at least some of the plurality of voxels which represent the volume within the anatomical structure based at least in part on the minimized divergence of the velocity field. Correcting velocity information for the at least some of the plurality of voxels may include correcting x velocity, y velocity and z velocity for the at least some of the plurality of voxels. Minimizing a divergence of a velocity field may include constructing a least squares divergence free approximation of the velocity field with a constraint of minimizing the divergence. Minimizing a divergence of a velocity field may include iteratively minimizing a residual divergence of the velocity field.

[0009] A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically correct phase aliasing may be summarized as including receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data including respective anatomical structure and velocity information for each of a plurality of voxels; analyzing, by the at least one processor, the set of MRI data to identify a time point which corresponds to a peak diastole of a cardiac cycle; analyzing, by the at least one processor, a variation of the velocity of at least some of the plurality of voxels over at least a portion of the cardiac cycle; and correcting, by the at least one processor, errors in the velocity of the at least some of the plurality of voxels based

at least in part on the analysis of the variation of the velocity of at least some of the plurality of voxels over at least a portion of a cardiac cycle.

[0010] Analyzing a variation of the velocity of at least some of the plurality of voxels and correcting errors in the velocity of at least some of the plurality of voxels may include, for each of the at least some of the plurality of voxels, tracking, by the at least one processor, a change in velocity between successive time points in the cardiac cycle; determining, by the at least one processor, that aliasing has occurred if the velocity varies by more than a magnitude of a velocity encoding parameter; responsive to determining that aliasing has occurred, incrementing, by the at least one processor, a wrap count if the velocity was reduced by more than the velocity encoding parameter; or decrementing, by the at least one processor, a wrap count if the velocity was increased by more than the velocity encoding parameter; and for each point in time, modifying, by the at least one processor, the velocity based at least in part on the accumulated wrap count for the voxel. Modifying the velocity based at least in part on the accumulated wrap count for the voxel may include increasing the velocity by an amount equal to product of the accumulated wrap count and two times the velocity encoding parameter.

[0011] The method may further include determining, by the at least one processor, whether the accumulated wrap count is zero over a cardiac cycle; and responsive to determining that the accumulated wrap count is not zero over a cardiac cycle, generating, by the at least one processor, an error signal.

[0012] The method may further include tracking, by the at least one processor, the number of voxels for which the accumulated wrap count is not zero over a cardiac cycle.

[0013] The method may further include correcting, by the at least one processor, errors in the velocity only for voxels for which the accumulated wrap count is not zero over a cardiac cycle.

[0014] The method may further include identifying, by the at least one processor, which of the plurality of voxels are likely to represent blood flow; and selecting, by the at least one processor, the at least some of the plurality of voxels for analysis based at least in part on the identification of which of the plurality of voxels are likely to represent blood flow.

[0015] A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically correct artifacts due to eddy currents may be summarized as including receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data including respective anatomical structure and velocity information for each of a plurality of voxels; receiving, by the at least one processor, an indication of which of the plurality of voxels represent static tissue; determining, by the at least one processor, at least one eddy current correction parameter based at least in part on the velocity information for the voxels which represent static tissue; and modifying, by the at least one processor, the velocity information for the plurality of voxels based at least in part the determined at least one eddy current correction parameter.

[0016] Receiving an indication of which of the plurality of voxels represent static tissue may include filtering, by the at least one processor, the MRI data to mask regions of air; filtering, by the at least one processor, the MRI data to mask regions of blood flow; and filtering, by the at least one

processor, the MRI data to mask regions of non-static tissue. Filtering the MRI data to mask regions of air may include masking regions with anatomy image values that are below a determined threshold.

[0017] The method may further include analyzing, by the at least one processor, a histogram for anatomy image values to determine the determined threshold.

[0018] The method may further include receiving, by the at least one processor, at least one filter parameter via a user interface communicatively coupled to the at least one processor.

[0019] A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically set an orientation of a multiplanar reconstruction may be summarized as including receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data including respective anatomical structure and velocity information for each of a plurality of voxels; presenting, by a display communicatively coupled to the at least one processor, the MRI data; receiving, by the at least one processor, selection of a central region of blood flow; determining, by the at least one processor, the direction of blood flow in the central region of blood flow; and responsive to determining the direction of blood flow, adjusting, by the at least one processor, the orientation of the multiplanar reconstruction so that the multiplanar reconstruction is on a plane that is perpendicular to the determine direction of blood flow. Determining the direction of blood flow in the central region of blood flow may include determining a time point which corresponds to peak blood flow; and determining the direction of blood flow at the time point which corresponds to peak blood flow.

[0020] A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically quantify blood flow in an anatomical volume may be summarized as including receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data including respective anatomical structure and velocity information for each of a plurality of voxels; identifying, by the at least one processor, an anatomical volume in the set of MRI data which may include a blood pool; identifying, by the at least one processor, a plane on a landmark that is at least approximately perpendicular to the flow of blood in the anatomical volume; generating, by the at least one processor, a contour defined by the intersection of the plane and the anatomical volume; and determining, by the at least one processor, the total flow of blood through the voxels in the generated contour. Determining the total flow of blood through the voxels in the generated contour may include, for each voxel in the generated contour, determining, by the at least one processor, a dot product of a normal vector of the plane and a velocity vector of the voxel; and summing, by the at least one processor, the determined dot products of each of the voxels in the generated contour.

[0021] The velocity information may include a plurality of velocity components, and determining the total flow of blood through the voxels in the generated contour may include for each velocity component, producing a velocity multiplanar reconstruction (MPR) in the plane of the generated contour, each of the velocity MPRs including a plurality of MPR pixels having pixel spacing at least similar to dimension of each of the voxels; for each MPR pixel inside the generated contour, determining, by the at least one

processor, a dot product of a normal vector of the plane and a velocity vector composed of the pixel values from the generated velocity MPRs; and summing, by the at least one processor, the determined dot products of each of the MPR pixels in the generated contour.

[0022] A method of operation for use with a medical imaging system to automatically determine a volume of a particular region may be summarized as including receiving, by at least one processor, a set of image data representative of an anatomical structure, the set of image data comprising respective anatomical structure information for each of a plurality of voxels; receiving, by the at least one processor, an indication of a primary axis of a volume of interest; generating, by the at least one processor, a plurality of spaced-apart slice planes along the primary axis of the volume of interest; for each of the plurality of slice planes, generating, by the at least one processor, a closed contour which defines a boundary of the volume of interest at the slice plane; generating, by the at least one processor, a three dimensional surface which connects the contours of all of the slice planes; and determining, by the at least one processor, the volume of the three dimensional surface. The set of image data may include respective anatomical structure information for each of a plurality of voxels at a plurality of time points, and the method may include, for each of the plurality of time points, generating, by the at least one processor, a plurality of spaced-apart slice planes along the primary axis of the volume of interest; for each of the plurality of slice planes, generating, by the at least one processor, a closed contour which defines a boundary of the volume of interest at the slice plane; generating, by the at least one processor, a three dimensional surface which connects the contours of all of the slice planes; and determining, by the at least one processor, the volume of the three dimensional surface. Receiving an indication of a primary axis may include receiving an indication of a primary axis that comprises one of a straight axis or a curved axis. The set of image data may include respective anatomical structure information for each of a plurality of voxels at a plurality of time points, and the primary axis may include a primary axis that moves over at least some of the plurality of time points. Receiving an indication of a primary axis may include receiving an indication of multiple axes.

[0023] A method of operation for use with a medical imaging system to automatically generate a connected mask may be summarized as including receiving, by at least one processor, a set of image data representative of an anatomical structure, the set of image data comprising respective anatomical structure information for each of a plurality of locations in space; identifying, by the at least one processor, a seed location; identifying, by the at least one processor, an intensity value for the seed location; and traversing, by the at least one processor, outward from the seed location to other locations to flood fill the connected mask according to flood fill criteria, the flood fill criteria including at least whether the intensity value of the each of the locations is within a specified threshold of the intensity value for the seed location.

[0024] The received set of image data may include three dimensional image data, and the method may further include generating, by the at least one processor, a two dimensional multiplanar reconstruction from the received image data. Moving outward from the seed location to other locations to flood fill the mask may include moving outward from the

seed location to other locations to flood fill the mask according to flood fill criteria which may include connectivity criteria. Moving outward from the seed location to other locations to flood fill the mask may include moving outward from the seed location to other locations to flood fill the mask according to flood fill criteria, the flood fill criteria including at least one of a radius constraint or a number of flood fill steps constraint.

[0025] The method may further include generating, by the at least one processor, a contour defined by a border of the connected mask; and generating, by the at least one processor, an approximation contour which is an approximation of the generated contour based at least in part on the generated contour. Generating an approximation contour may include generating an approximation contour using a spline which places control points at areas of high curvature.

[0026] The method may further include dilating, by the at least one processor, the approximation contour by a determined dilation factor to generate a dilated approximation contour.

[0027] A method of operation for use with a medical imaging system may be summarized as including receiving, by at least one processor, a set of image data representative of an anatomical structure, the set of image data comprising respective anatomical structure information for each of a plurality of locations in three dimensional space; tracking, by the at least one processor, the position and orientation of the anatomical structure at a plurality of time points; at each of the plurality of time points, generating, by the at least one processor, a contour for the anatomical structure; and determining, by the at least one processor, a measurement associated with the anatomical structure using the generated contours. Determining a measurement associated with the anatomical structure using the generated contours may include accounting for movement of the anatomical structure over at least some of the plurality of time points. Accounting for movement of the anatomical structure may include accounting for the linear and angular velocity of the anatomical structure over the at least some of the plurality of time points. Determining a measurement associated with the anatomical structure using the generated contours may include determining flow through the anatomical structure.

[0028] The method may further include presenting, by the at least one processor, a multiplanar reconstruction of the anatomical structure on a display at a plurality of time points, wherein the position and orientation of the multiplanar reconstruction tracks the generated contours over the plurality of time points.

[0029] The method may further include presenting, by the at least one processor, a multiplanar reconstruction of the anatomical structure on a display; and presenting, by the at least one processor, the generated contours as semi-transparent if the multiplanar reconstruction is out of plane with the generated contours.

[0030] A processor-based device may have at least one processor and at least one nontransitory processor-readable medium communicatively coupled to the at least one processor, and operable to perform any of the above methods.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0031] In the drawings, identical reference numbers identify similar elements or acts. The sizes and relative positions of elements in the drawings are not necessarily drawn to

scale. For example, the shapes of various elements and angles are not necessarily drawn to scale, and some of these elements may be arbitrarily enlarged and positioned to improve drawing legibility. Further, the particular shapes of the elements as drawn, are not necessarily intended to convey any information regarding the actual shape of the particular elements, and may have been solely selected for ease of recognition in the drawings.

[0032] FIG. 1 is a schematic view of a networked environment including at least one MRI acquisition system and at least one image processing system, the MRI acquisition system located in a clinical setting and the image processing system located remotely from the MRI acquisition system and communicatively coupled therewith over one or more networks, according to one illustrated embodiment.

[0033] FIG. 2 is a functional block diagram of an MRI acquisition system and an MRI image processing and analysis system that provides MRI image processing and analysis services, according to one illustrated embodiment.

[0034] FIGS. 3A-3B is a flow diagram of an example push process executable by at least one processor, according to one illustrated embodiment.

[0035] FIGS. 4A-4B is a flow diagram of an example process of monitoring for artifacts and arching executable by at least one processor, according to one illustrated embodiment.

[0036] FIG. 5 is a schematic illustration of a PHI service pipeline, according to one illustrated embodiment.

DETAILED DESCRIPTION

[0037] In the following description, certain specific details are set forth in order to provide a thorough understanding of various disclosed embodiments. However, one skilled in the relevant art will recognize that embodiments may be practiced without one or more of these specific details, or with other methods, components, materials, etc. In other instances, well-known structures associated with MRI machines, computer systems, server computers, and/or communications networks have not been shown or described in detail to avoid unnecessarily obscuring descriptions of the embodiments.

[0038] Unless the context requires otherwise, throughout the specification and claims which follow, the word “comprise” and variations thereof, such as, “comprises” and “comprising” are synonymous with “including,” and are inclusive or open-ended (i.e., does not exclude additional, unrecited elements or method acts).

[0039] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0040] As used in this specification and the appended claims, the singular forms “a,” “an,” and “the” include plural referents unless the content clearly dictates otherwise. It should also be noted that the term “or” is generally employed in its sense including “and/or” unless the content clearly dictates otherwise.

[0041] The headings and Abstract of the Disclosure provided herein are for convenience only and do not interpret the scope or meaning of the embodiments.

[0042] Many of the implementations described herein take advantage of a 4-D flow MRI data set, which essentially captures MRI magnitude and phase information for a three dimensional volume over a period of time. This approach may allow capture or acquisition of MRI data sets without requiring breath holding or synchronization or gating to a patient’s cardiac or pulmonary cycles. Instead, MRI data sets are captured or acquired, and imaging processing and analysis employed to derive the desired information, for example by re-binning acquired information based on the cardiac and pulmonary cycles. This essentially pushes what is normally time-intensive acquisition operations to the imaging processing and analysis stage. As way of a simplified analogy, in some respects such may be thought of as capturing a movie of the anatomical structure (e.g., chest, heart) without concern over a patient’s pulmonary or cardiac cycles, the processing the captured movie to account for relative movement introduced by the pulmonary and cardiac cycles. The captured information includes both magnitude information, which is indicative of anatomical structure, and phase information which is indicative of velocity. The phase information allows distinction between static and non-static tissue, for example allowing non-static tissue (e.g., blood, air) to be distinguished from static tissue (e.g., fat, bone). The phase information also allows certain non-static tissue (e.g., air) to be distinguished from other non-static tissue (e.g., blood). This may advantageously allow automated or even autonomous segmentation between tissues, and/or distinguishing atrial blood flow from venous blood flow. This may advantageously allow automated or even autonomous generation of flow visualization information, which may be superimposed on anatomical information. This may also advantageously allow automated or even autonomous flow quantification, identifying abnormalities and/or verifying results.

[0043] The workflow may general divided into three portions, sequentially: 1) image acquisition, 2) image reconstruction, and 3) image processing or post-processing and analysis. Alternatively, the workflow may be divided into 1) operational, 2) preprocessing, and 3) visualization and quantification.

[0044] Image acquisition may include determining, defining, generating or otherwise setting one or more pulse sequences, which are used to run the MRI machine (e.g., control magnets) and acquire raw MRI. Use of a 4D flow pulse sequence allows capture of not only anatomical structure, which is represented by magnitude, but of velocity, which is represented by phase. At least one of the methods or techniques described herein, generation of patient specific 4D pulse sequences, occurs during or as part of image acquisition portion. Image reconstruction may, for example, employ fast Fourier transformations, and result in MRI data sets, often in a form compatible with the DICOM standard. Image reconstruction has traditional been computationally intensive often relying on supercomputers. The requirement for such is a significant burden to many clinical facilities. Many of the methods and techniques described herein occur during or as part of the imaging processor or post-processing and analysis. Such can include error detection and/or error correction, segmentation, visualization including fusion of flow related information and images of anatomical struc-

tures, quantification, identification of abnormalities including shunts, verification including identification of spurious data. Alternatively, error detection and/or error correction may occur during the preprocessing portion.

[0045] FIG. 1 shows a networked environment 100 according to one illustrated embodiment, in which one or more MRI acquisition system (one shown) 102 are communicatively coupled to at least one image processing and analysis system 104 via one or more networks 106a, 106b (two shown, collectively 106).

[0046] The MRI acquisition system 102 is typically located at a clinical facility, for instance a hospital or dedicated medical imaging center. Various techniques and structures, as explained herein, may advantageously allow the image processing and analysis system 104 to be remotely located from the MRI acquisition system 102. The image processing and analysis system 104 may, for example, be located in another building, city, state, province or even country.

[0047] The MRI acquisition system 102 may, for example, include an MRI machine 108, a computer system 110 and an MRI operator's system 112. The MRI machine 108 may include a main magnet 114, which is typically an annular array of coils having a central or longitudinal bore 116. The main magnet 108 is capable of producing a strong stable magnetic field (e.g., 0.5 Tesla to 3.0 Tesla). The bore 116 is sized to receive at least a portion of an object to be imaged, for instance a human body 118. When used in medical imaging applications, the MRI machine 108 typically includes a patient table 120 which allows a prone patient 118 to be easily slid or rolled into and out of the bore 116.

[0048] The MRI machine also includes a set of gradient magnets 122 (only one called out). The gradient magnets 122 produce a variable magnetic field that is relatively smaller than that produced by the main magnet 114 (e.g., 180 Gauss to 270 Gauss), allowing selected portions of an object (e.g., patient) to be imaged. MRI machine 108 also include radio frequency (RF) coils 124 (only one called out) which are operated to apply radiofrequency energy to selected portions of the object (e.g., patient 118) to be imaged. Different RF coils 124 may be used for imaging different structures (e.g., anatomic structures). For example, one set of RF coils 124 may be appropriate for imaging a neck of a patient, while another set of RF coils 124 may be appropriate for imaging a chest or heart of the patient. MRI machines 108 commonly include additional magnets, for example resistive magnets and/or permanent magnets.

[0049] The MRI machine 108 typically includes, or is communicatively coupled to, a processor-based MRI control system 126 used to control the magnets and/or coils 114, 122, 124. The processor-based control system 126 may include one or more processors, non-transitory computer- or processor-readable memory, drive circuitry and/or interface components to interface with the MRI machine 108. The processor-based control system 126 may, in some implementations, also perform some preprocessing on data resulting from the MRI operation.

[0050] An MRI operator's system 128 may include a computer system 130, monitor or display 132, keypad and/or keyboard 134, and/or a cursor control device such as a mouse 136, joystick, trackpad, trackball or the like. The MRI operator's system 128 may include or read computer- or processor executable instructions from one or more non-transitory computer- or processor-readable medium, for

instance spinning media 138 such as a magnetic or optical disk. The operator's system 128 may allow a technician to operate the MRI machine 108 to capture MRI data from a patient 118. Various techniques, structures and features described herein may allow MRI machine 108 operation by a technician without requiring the presence of a clinician or physician. Such may advantageously significantly lower costs of MRI procedures. Also as described herein, various techniques, structures and features may allow MRI procedures to be performed much more quickly than using conventional techniques. Such may advantageously allow higher throughput for each MRI installation, amortizing cost of the capital intensive equipment over a much larger number of procedures. For example, high computational power computers may be located remotely from the clinical setting, and may be used to server multiple clinical facilities. The various techniques, structures and features described herein may also additionally or alternatively advantageously reduce the time that each patient is exposed to the MRI procedure, reducing or alleviating the anxiety that often accompanies undergoing an MRI procedure. For instance, eliminating the need for breath holding and/or synchronizing with a patient's pulmonary and/or cardiac cycles via image processing and analysis techniques described herein may significantly reduce the time for acquisition, for example to eight to ten minutes.

[0051] The image processing and analysis system 104 may include one or more servers 139 to handle incoming requests and responses, and one or more rendering or image processing and analysis computers 140. The server(s) 139 may, for example take the form of one or more server computers, workstation computers, supercomputers, or personal computers, executing server software or instructions. The one or more rendering or image processing and analysis computers 140 may take the form of one or more computers, workstation computers, supercomputers, or personal computers, executing image processing and/or analysis software or instructions. The one or more rendering or image processing and analysis computers 140 will typically employ one, and preferably multiple, graphical processing units (GPUs) or GPU cores.

[0052] The image processing and analysis system 104 may include one or more non-transitory computer-readable medium 142 (e.g., magnetic or optical hard drives, RAID, RAM, Flash) that stores processor-executable instructions and/or data or other information. The image processing and analysis system 104 may include one or more may include one or more image processing and analysis operator's systems 144. The image processing and analysis operator's system 144 may include a computer system 146, monitor or display 148, keypad and/or keyboard 150, and/or a cursor control device such as a mouse 152, joystick, trackpad, trackball or the like. The image processing and analysis operator's system 144 may be communicatively coupled to the rendering or image processing and analysis computer(s) 140 via one or more networks, for instance a LAN 154. While many image processing techniques and analysis may be fully automated, the image processing and analysis operator's system may allow a technician to perform certain image processing and/or analysis operations on MRI data captured from a patient.

[0053] While illustrated as a single nontransitory computer- or processor-readable storage medium 142, in many implementations the nontransitory computer- or processor-

readable storage medium **142** may constitute a plurality of nontransitory storage media. The plurality of nontransitory storage media may be commonly located at a common location, or distributed at a variety of remote locations. Thus, a database of raw MRI data, preprocessed MRI data and/or processed MRI data may be implemented in one, or across more than one, nontransitory computer- or processor-readable storage media. Such database(s) may be stored separately from one another on separate computer- or processor-readable storage medium **142** or may be stored on the same computer- or processor-readable storage medium **142** as one another. The computer- or processor-readable storage medium **142** may be co-located with the image processing and analysis system **104**, for example, in the same room, building or facility. Alternatively, the computer- or processor-readable storage medium **142** may be located remotely from the image processing and analysis system **104**, for example, in a different facility, city, state or country. Electronic or digital information, files or records or other collections of information may be stored at specific locations in non-transitory computer- or processor-readable media **142**, thus are logically addressable portions of such media, which may or may not be contiguous.

[**0054**] As noted above, the image processing and analysis system **104** may be remotely located from the MRI acquisition system **102**. The MRI acquisition system **102** and the image processing and analysis system **104** are capable of communications, for example via one or more communications channels, for example local area networks (LANs) **106a** and Wide Area Networks (WANs) **106b**. The networks **106** may, for instance include packet switched communications networks, such as the Internet, Worldwide Web portion of the Internet, extranets, and/or intranets. The networks **106** may take the form of various other types of telecommunications networks, such as cellular phone and data networks, and plain old telephone system (POTS) networks. The type of communications infrastructure should not be considered limiting.

[**0055**] As illustrated in FIG. 1, the MRI acquisition system **102** is communicatively coupled to the first LAN **106a**. The first LAN **106a** may be a network operated by or for the clinical facility, providing local area communications for the clinical facility. The first LAN **106a** is communicatively coupled to the WAN (e.g., Internet) **106b**. A first firewall **156a** may provide security for the first LAN.

[**0056**] Also as illustrated in FIG. 1, the image processing and analysis system **104** is communicatively coupled to the second LAN **154**. The second LAN **154** may be a network operated by or for an image processing facility or entity, providing local area communications for the image processing facility or entity. The second LAN **154** is communicatively coupled to the WAN **106b** (e.g., Internet). A second firewall **156b** may provide security for the second LAN **154**.

[**0057**] The image processing facility or entity may be independent from the clinical facility, for example an independent business providing services to one, two or many clinical facilities.

[**0058**] While not illustrated, the communications network may include one or more additional networking devices. The networking devices may take any of a large variety of forms, including servers, routers, network switches, bridges, and/or modems (e.g., DSL modem, cable modem), etc.

[**0059**] While FIG. 1 illustrates a representative networked environment **100**, typical networked environments may

include many additional MRI acquisition systems, image processing and analysis system **104**, computer systems, and/or entities. The concepts taught herein may be employed in a similar fashion with more populated networked environments than that illustrated. For example, a single entity may provide image processing and analysis services to multiple diagnostic entities. One or more of the diagnostic entities may operate two or more MRI acquisition systems **102**. For example, a large hospital or dedicated medical imaging center may operate two, three or even more MRI acquisition systems at a single facility. Typically, the entity that provides the image processing and analysis services will operate multiple entity may provide image processing and analysis systems **104** which may include two, three or even hundreds of rendering or image processing and analysis computers **140**.

[**0060**] FIG. 2 shows a networked environment **200** comprising one or more image processing and analysis systems **104** (only one illustrated) and one or more associated nontransitory computer- or processor-readable storage medium **204** (only one illustrated). The associated nontransitory computer- or processor-readable storage medium **204** is communicatively coupled to the image processing and analysis system(s) **104** via one or more communications channels, for example, one or more parallel cables, serial cables, or wireless channels capable of high speed communications, for instance, via FireWire®, Universal Serial Bus® (USB) 2 or 3, and/or Thunderbolt®, Gigabyte Ethernet®.

[**0061**] The networked environment **200** also comprises one or more end MRI acquisition systems **102** (only one illustrated). The MRI acquisition system(s) **102** are communicatively coupled to the image processing and analysis system(s) **104** by one or more communications channels, for example, one or more wide area networks (WANs) **210**, for instance the Internet or Worldwide Web portion thereof.

[**0062**] In operation, the MRI acquisition systems **102** typically function as a client to the image processing and analysis system **104**. In operation, the image processing and analysis systems **104** typically functions as a server to receive requests or information (e.g., MRI data sets) from the MRI acquisition systems **102**. Described herein is an overall process which employs an asynchronous command and imaging pipeline that allows the image processing and analysis to be performed remotely (e.g., over a WAN) from the MRI acquisition system **102**. This approach provides for a number of distinctive advantages, for instance allowing the MRI acquisition system(s) **102** to be operated by a technician without requiring the presence of a clinician (e.g., physician). Various techniques or approaches are also described to enhance security, while allowing access to medical imaging data as well as private patient specific health information.

[**0063**] While illustrated as located remotely from the MRI acquisition system(s) **102**, in some implementations the image processing and analysis systems **104** may be co-located with the MRI acquisition system **102**. In other implementations, one or more of the operations or functions described herein may be performed by the MRI acquisition system **102** or via a processor-based device co-located with the MRI acquisition system **102**.

[**0064**] The image processing and analysis systems **104** receive MRI data sets, perform image processing on the MRI data sets, and provide the processed MRI data sets, for

example to a clinician for review. The image processing and analysis systems **104** may, for example, perform error detection and/or correction on MRI data sets, for example phase error correction, phase aliasing detection, signal unwrapping, and/or detection and/or correction of various artifacts. Phase error is related to phase, as is phase aliasing. Signal unwrapping is related to magnitude. Various other artifacts may be related to phase and/or magnitude.

[0065] The image processing and analysis systems **104** may, for example, perform segmentation, distinguishing between various tissue type. The image processing and analysis systems **104** may, for example, perform quantification, for instance comparing blood flow into and out of a closed anatomical structure or through two or more anatomical structures. The image processing and analysis systems **104** may advantageously use quantification to verify results, for example confirming identification of a certain tissue and/or providing an indication of an amount of certainty in the results. Additionally, the image processing and analysis systems **104** may advantageously use quantification to identify the existence of a shunt.

[0066] In some implementations, the image processing and analysis systems **104** may generate images which reflect blood flow, for example including distinguishing between arterial and venous blood flow. For instance, the image processing and analysis systems **104** may employ a first color map (e.g., blue) to indicate arterial blood flow and a second color map (e.g., red) to indicate venous blood flow. The image processing and analysis systems **104** may indicate aberrations (e.g., shunt) using some other, distinctive color or visual emphasis. Numerous different techniques are described for distinguishing between different tissues as well as between arterial and venous blood flow. Flow visualization may be superimposed, for instance as one or more layers, on or over visual representations of anatomical structure or magnitude data.

[0067] In some implementations, the image processing and analysis systems **104** may generate a patient specific 4D-flow protocol for use in operating an MRI acquisition system **102** with a specific patient. Such may include setting an appropriate velocity encoding (VENC) for operation of the MRI machine.

[0068] The image processing and analysis systems **104** may perform one or more of these operations or functions autonomously, without human input. Alternatively, the image processing and analysis systems **104** may perform one or more of these operations or functions based on human input, for example human input which identifies a point, location or plane, or which otherwise identifies a characteristic of anatomical tissue. Some planes and/or views may be predefined, allowing the operator, user or clinician to simply select a plane (e.g., a valve plane) or a denominated view (e.g., 2 chamber view, 3 chamber view, 4 chamber view) to quickly and easily obtain the desired view.

[0069] The networked environment **200** may employ other computer systems and network equipment, for example, additional servers, proxy servers, firewalls, routers and/or bridges. The image processing and analysis systems **104** will at times be referred to in the singular herein, but this is not intended to limit the embodiments to a single device since in typical embodiments there may be more than one image processing and analysis systems **104** involved. Unless described otherwise, the construction and operation of the various blocks shown in FIG. 2 are of conventional design.

As a result, such blocks need not be described in further detail herein, as they will be understood by those skilled in the relevant art.

[0070] The image processing and analysis systems **104** may include one or more processing units **212a**, **212b** (collectively **212**), a system memory **214** and a system bus **216** that couples various system components, including the system memory **214** to the processing units **212**. The processing units **212** may be any logic processing unit, such as one or more central processing units (CPUs) **212a**, digital signal processors (DSPs) **212b**, application-specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), etc. The system bus **216** can employ any known bus structures or architectures, including a memory bus with memory controller, a peripheral bus, and/or a local bus. The system memory **214** includes read-only memory (“ROM”) **218** and random access memory (“RAM”) **220**. A basic input/output system (“BIOS”) **222**, which can form part of the ROM **218**, contains basic routines that help transfer information between elements within the image processing and analysis system(s) **104**, such as during start-up.

[0071] The image processing and analysis system(s) **104** may include a hard disk drive **224** for reading from and writing to a hard disk **226**, an optical disk drive **228** for reading from and writing to removable optical disks **232**, and/or a magnetic disk drive **230** for reading from and writing to magnetic disks **234**. The optical disk **232** can be a CD-ROM, while the magnetic disk **234** can be a magnetic floppy disk or diskette. The hard disk drive **224**, optical disk drive **228** and magnetic disk drive **230** may communicate with the processing unit **212** via the system bus **216**. The hard disk drive **224**, optical disk drive **228** and magnetic disk drive **230** may include interfaces or controllers (not shown) coupled between such drives and the system bus **216**, as is known by those skilled in the relevant art. The drives **224**, **228** and **230**, and their associated computer-readable media **226**, **232**, **234**, provide nonvolatile storage of computer-readable instructions, data structures, program modules and other data for the image processing and analysis system(s) **104**. Although the depicted image processing and analysis systems **104** is illustrated employing a hard disk **224**, optical disk **228** and magnetic disk **230**, those skilled in the relevant art will appreciate that other types of computer-readable media that can store data accessible by a computer may be employed, such as WORM drives, RAID drives, magnetic cassettes, flash memory cards, digital video disks (“DVD”), Bernoulli cartridges, RAMs, ROMs, smart cards, etc.

[0072] Program modules can be stored in the system memory **214**, such as an operating system **236**, one or more application programs **238**, other programs or modules **240** and program data **242**. Application programs **238** may include instructions that cause the processor(s) **212** to perform image processing and analysis on MRI image data sets. For example, the application programs **238** may include instructions that cause the processor(s) **212** to perform phase error correction on phase or velocity related data. For example, the application programs **238** may include instructions that cause the processor(s) **212** to correct for phase aliasing. Also for example, the application programs **238** may include instructions that cause the processor(s) **212** to perform signal unwrapping. Additionally or alternatively, the application programs **238** may include instructions that cause the processor(s) **212** to identify and/or correct for artifacts.

[0073] The application programs 238 may include instructions that cause the processor(s) 212 to, for example, perform segmentation, distinguishing between various tissue type. The application programs 238 may include instructions that cause the processor(s) 212 to perform quantification, for instance comparing blood flow into and out of a closed anatomical structure or through two or more anatomical structures. The application programs 238 may include instructions that cause the processor(s) 212 to use quantification to verify results, for example confirming identification of a certain tissue and/or providing an indication of an amount of certainty in the results. The application programs 238 may include instructions that cause the processor(s) 212 to use quantification to identify the existence of a shunt.

[0074] The application programs 238 may include instructions that cause the processor(s) 212 to generate images which reflect blood flow, for example distinguishing between arterial and venous blood flow. For instance, a first color map (e.g., blue) may be used to indicate arterial blood flow and a second color map (e.g., red) to indicate venous blood flow. Aberrations (e.g., shunt) may be indicated using some other, distinctive color or visual emphasis. Color transfer functions may be applied to generate the color maps. The application programs 238 may include instructions that cause the processor(s) 212 to superimpose visualization of flow (e.g., MRI phase data indicative of blood flow velocity and/or volume) on visualization or rendered images of anatomy (e.g., MRI magnitude data). The instructions may cause the flow visualization to be rendered as one or more layers on the images of the anatomy to provide a fusion of anatomy (i.e., magnitude) and flow (i.e., phase) information, for example as a color heat map and/or as vectors (e.g. arrow icons) with direction and magnitude (e.g., represented by length, line weight). The instructions may additionally or alternatively cause the generation of spatial mappings or visualization of signal dispersion, turbulence and/or pressure, which may be overlaid or superimposed on a spatial mapping or visualization of anatomical structure. Fusing visualization of phase or velocity related information with visualization of anatomical information or visual representations of anatomical structures may facilitate the identification of anatomical landmarks. The instructions may make use of sets or arrays of graphics processing units or GPUs to quickly render the visualizations.

[0075] Transfer functions may also be applied to determine which visual effects (e.g., color) to apply to which tissue. For example, arterial blood flow may be colored in shades of blue and venous blood flow in shades of red, while fat tissue colored as yellow. Anatomical structure, represented as magnitude in the MRI image data set, may, for example, be visualized using grey scale. Depth of view may be operator or user adjustable, for example via a slider control on a graphical user interface. Thus, visualization may be in the form a fusion view that advantageously fuses a visual representation of velocity information with a visual representation of anatomical information or representation.

[0076] The application programs 238 may include instructions that cause the processor(s) 212 to generate a patient specific 4D-flow protocol for use in operating an MRI acquisition system 102 with a specific patient. Such may be based on patient specific input, for example provided by a technician, and may be based on the particular MRI machine being used to capture the MRI data set.

[0077] The application programs 238 may include instructions that cause the processor(s) 212 to receive image data sets from the MRI acquisition system, process and/or analyze the image data sets, and provide processed and/or analyzed images and other information to users remotely located from the image processing, in a time sensitive and secure manner. Such is described in detail herein with reference to the various Figures.

[0078] The system memory 214 may also include communications programs, for example, a server 244 that causes the image processing and analysis system(s) 104 to serve electronic information or files via the Internet, intranets, extranets, telecommunications networks, or other networks as described below. The server 244 in the depicted embodiment is markup language based, such as Hypertext Markup Language (HTML), Extensible Markup Language (XML) or Wireless Markup Language (WML), and operates with markup languages that use syntactically delimited characters added to the data of a document to represent the structure of the document. A number of suitable servers may be commercially available such as those from Mozilla, Google, Microsoft and Apple Computer.

[0079] While shown in FIG. 2 as being stored in the system memory 214, the operating system 236, application programs 238, other programs/modules 240, program data 242 and server 244 can be stored on the hard disk 226 of the hard disk drive 224, the optical disk 232 of the optical disk drive 228 and/or the magnetic disk 234 of the magnetic disk drive 230.

[0080] An operator can enter commands and information into the image processing and analysis system(s) 104 through input devices such as a touch screen or keyboard 246 and/or a pointing device such as a mouse 248, and/or via a graphical user interface. Other input devices can include a microphone, joystick, game pad, tablet, scanner, etc. These and other input devices are connected to one or more of the processing units 212 through an interface 250 such as a serial port interface that couples to the system bus 216, although other interfaces such as a parallel port, a game port or a wireless interface or a universal serial bus ("USB") can be used. A monitor 252 or other display device is coupled to the system bus 216 via a video interface 254, such as a video adapter. The image processing and analysis system(s) 104 can include other output devices, such as speakers, printers, etc.

[0081] The image processing and analysis systems 104 can operate in a networked environment 200 using logical connections to one or more remote computers and/or devices. For example, the image processing and analysis 104 can operate in a networked environment 200 using logical connections to one or more MRI acquisition systems 102. Communications may be via a wired and/or wireless network architecture, for instance, wired and wireless enterprise-wide computer networks, intranets, extranets, and/or the Internet. Other embodiments may include other types of communications networks including telecommunications networks, cellular networks, paging networks, and other mobile networks. There may be any variety of computers, switching devices, routers, bridges, firewalls and other devices in the communications paths between the image processing and analysis systems 104, the MRI acquisition systems 102.

[0082] The MRI acquisition systems 102 will typically take the form of an MRI machine 108 and one or more

associated processor-based devices, for instance an MRI control system 126 and/or MRI operator's system 128. The MRI acquisition systems 102 capture MRI information or data sets from patients. Thus, in some instances the MRI acquisition systems 102 may be denominated as front end MRI acquisition systems or MRI capture systems, to distinguish such from the MRI image processing and analysis system(s) 104, which in some instances may be denominated as MRI backend systems. The MRI acquisition systems 102 will at times each be referred to in the singular herein, but this is not intended to limit the embodiments to a single MRI acquisition system 102. In typical embodiments, there may be more than one MRI acquisition system 102 and there will likely be a large number of MRI acquisition systems 102 in the networked environment 200.

[0083] The MRI acquisition systems 102 may be communicatively coupled to one or more server computers (not shown). For instance, MRI acquisition systems 102 may be communicatively coupled via one or more diagnostic facility server computers (not shown), routers (not shown), bridges (not shown), LANs 106a (FIG. 1), etc., which may include or implement a firewall 156a (FIG. 1). The server computers (not shown) may execute a set of server instructions to function as a server for a number of MRI acquisition systems 102 (i.e., clients) communicatively coupled via a LAN 106a at a clinical facility or site, and thus act as intermediaries between the MRI acquisition systems 102 and the MRI image processing and analysis system(s) 104. The MRI acquisition systems 102 may execute a set of client instructions to function as a client of the server computer(s), which are communicatively coupled via a WAN.

[0084] The MRI control system 126 typically includes one or more processor (e.g., microprocessors, central processing units, digital signal processors, graphical processing units) and non-transitory processor-readable memory (e.g., ROM, RAM, Flash, magnetic and/or optical disks). The MRI operator's system 128 may take the form of a computer, for instance personal computers (e.g., desktop or laptop computers), net book computers, tablet computers, smart phones, personal digital assistants, workstation computers and/or mainframe computers, and the like, executing appropriate instructions.

[0085] The MRI operator's system 128 may include one or more processing units 268, system memories 269 and a system bus (not shown) that couples various system components including the system memory 269 to the processing unit 268.

[0086] The processing unit 268 may be any logic processing unit, such as one or more central processing units (CPUs), digital signal processors (DSPs), application-specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), graphical processing units (GPUs), etc. Non-limiting examples of commercially available computer systems include, but are not limited to, an 80x86 or Pentium series microprocessor from Intel Corporation, U.S.A., a PowerPC microprocessor from IBM, a Sparc microprocessor from Sun Microsystems, Inc., a PA-MSC series microprocessor from Hewlett-Packard Company, a 68xxx series microprocessor from Motorola Corporation, an ATOM processor, or an A4 or A5 processor. Unless described otherwise, the construction and operation of the various blocks of the MRI acquisition systems 102 shown in FIG. 2 are of conventional design. As a result, such blocks need not be

described in further detail herein, as they will be understood by those skilled in the relevant art.

[0087] The system bus can employ any known bus structures or architectures, including a memory bus with memory controller, a peripheral bus, and a local bus. The system memory 269 includes read-only memory ("ROM") 270 and random access memory ("RAM") 272. A basic input/output system ("BIOS") 271, which can form part of the ROM 270, contains basic routines that help transfer information between elements within the MRI acquisition systems 102, such as during start-up.

[0088] The MRI operator's system 128 may also include one or more media drives 273, e.g., a hard disk drive, magnetic disk drive, WORM drive, and/or optical disk drive, for reading from and writing to computer-readable storage media 274, e.g., hard disk, optical disks, and/or magnetic disks. The nontransitory computer-readable storage media 274 may, for example, take the form of removable media. For example, hard disks may take the form of a Winchester drive, and optical disks can take the form of CD-ROMs, while magnetic disks can take the form of magnetic floppy disks or diskettes. The media drive(s) 273 communicate with the processing unit 268 via one or more system buses. The media drives 273 may include interfaces or controllers (not shown) coupled between such drives and the system bus, as is known by those skilled in the relevant art. The media drives 273, and their associated nontransitory computer-readable storage media 274, provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for MRI acquisition system(s) 102. Although described as employing computer-readable storage media 274 such as hard disks, optical disks and magnetic disks, those skilled in the relevant art will appreciate that MRI operator's system(s) 128 may employ other types of nontransitory computer-readable storage media that can store data accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks ("DVD"), Bernoulli cartridges, RAMs, ROMs, smart cards, etc. Data or information, for example, electronic or digital files or data or metadata related to such can be stored in the nontransitory computer-readable storage media 274.

[0089] Program modules, such as an operating system, one or more application programs, other programs or modules and program data, can be stored in the system memory 269. Program modules may include instructions for accessing a Website, extranet site or other site or services (e.g., Web services) and associated Webpages, other pages, screens or services hosted or provided by the MRI processing and analysis system(s) 104.

[0090] In particular, the system memory 269 may include communications programs that permit the MRI acquisition system(s) 102 to exchange electronic or digital information or files or data or metadata with the MRI image processing and/or analysis services provided by the MRI processing and analysis system(s) 104. The communications programs may, for example, be a Web client or browser that permits the MRI acquisition system(s) 102 to access and exchange information, files, data and/or metadata with sources such as Web sites of the Internet, corporate intranets, extranets, or other networks. Such may require that an end user client have sufficient right, permission, privilege or authority for accessing a given Website, for example, one hosted by the MRI processing and analysis system(s) 104. As discussed herein, patient identifying data may reside on systems oper-

ated by or for the clinical facility, and may not be accessible by or through the systems operated by or for the image processing facility or the image processing facility personnel. The browser may, for example, be markup language based, such as Hypertext Markup Language (HTML), Extensible Markup Language (XML) or Wireless Markup Language (WML), and may operate with markup languages that use syntactically delimited characters added to the data of a document to represent the structure of the document.

[0091] While described as being stored in the system memory 269, the operating system, application programs, other programs/modules, program data and/or browser can be stored on the computer-readable storage media 274 of the media drive(s) 273. An operator can enter commands and information into the MRI operator's system(s) 128 via a user interface 275 through input devices such as a touch screen or keyboard 276 and/or a pointing device 277 such as a mouse. Other input devices can include a microphone, joystick, game pad, tablet, scanner, etc. These and other input devices are connected to the processing unit 269 through an interface such as a serial port interface that couples to the system bus, although other interfaces such as a parallel port, a game port or a wireless interface or a universal serial bus ("USB") can be used. A display or monitor 278 may be coupled to the system bus via a video interface, such as a video adapter. The MRI operator system (s) 128 can include other output devices, such as speakers, printers, etc.

[0092] The MRI image processing and analysis system may build a static interface, which allows various tissue types to be subtracted or added to an MRI 4D flow data set. For example, static tissues such as fat or bone may be distinguished from non-static tissues such as air or flowing blood. The MRI image processing and analysis system may further autonomously distinguish between various non-static tissues, for instance distinguishing between air (e.g., lungs) and flowing blood. Further, the MRI image processing and analysis system may distinguish between arterial and venous blood flows.

[0093] For instance, the MRI image processing and analysis system may employ fast Fourier transformation to identify blood tissue, which is expected to have a pulsatile pattern or waveform. Air or lung will tend to have a random appear pattern over a defined volume, as velocity of neighboring voxels are compared. For instance, voxels with strong or fast velocities are typically indicative of air. The MRI data sets may be rather large, for example $256 \times 256 \times 256 \times 20$ time points. The MRI image processing and analysis system may rely on gradients (e.g., gradient descent method) to detect different tissue types, and may advantageously employ a numerical approach rather than an analytic solution approach to quickly handle the relatively large MRI data sets. By controlling the number of significant digits (e.g., 2) of the numerical approach, the MRI image processing and analysis system may achieve very fast (e.g., 1 second as opposed to 30 minutes) results, while still obtaining results that are sufficiently accurate for the particular application.

[0094] In some implementations, different tissue types may be subtracted from the patient MRI data set, one at a time. For example, subtracting air or lung, subtracting blood, separating atrial from venous flow, subtracting bone, leaving fat. Notably, fat is static, so each voxel representing fat should have zero velocity associated therewith. The MRI

image processing and analysis system may advantageously employ such a ground truth to correct MRI data set for all tissue types.

[0095] If a non-zero velocity is found for fat type tissue, this can be used to adjust the entire set of data (e.g., for all tissue). For example, the MRI image processing and analysis system may generate or create a polynomial model based on an identified area or volume (e.g., fat or soft tissue). Such may be a simple polynomial (e.g., ax^2+bx+c) or a much more complex polynomial (e.g., non-rational uniform b-spline). The MRI image processing and analysis system may find the coefficients to the polynomial fits the image, for example using linear regression techniques or linear algebra techniques. This results in a model which the MRI image processing and analysis system may apply to (e.g., subtract from) the whole field, not just the fat or soft tissue.

[0096] In one implementations, a replica body is imaged to create a reference set of data or "phantom" model which can be subtracted from actually patient data. The replica body may be formed of materials that mimic the MRI response of an actually body, although will not have blood flow. A phase gradient in reference set of data or "phantom" model may represent noise (e.g., random noise), and can be used to correct a phase shift. This approach advantageously avoids the need to generate a polynomial fit to the 3D data. The generated reference set or phantom model may be valid over a number of months of MRI machine operation, although a new set of reference data or phantom model should be generated if the MRI machine is serviced or moved.

[0097] The MRI image processing and analysis system may define various filters or mask for removing different tissue types or for removing either venous or atrial blood flow. Filters or masks may remove anomalous blood flow, such as blood flow outside some reasonable range (e.g., too high or fast, too slow or low) or where blood appears to be flowing in an anatomical structure (e.g., bone) where there should be no blood flow. A filter or mask may also be defined to display only voxels having magnitudes with an absolute value greater than some threshold value. A filter or mask may also be defined to display only voxels with an absolute value of the cross product of magnitude and a velocity vector which absolute value is greater than some defined threshold. Further a filter or mask may be defined that shows only voxels having vectors in a same direction as the vectors of neighboring voxels, to for instance identify or view high velocity jets. Notably, velocity vectors of neighboring voxels are in different directions may be an indication of noise.

Pre-Processing

[0098] Mass Conservation Correction Error Reduction

[0099] The goal of this pre-processing algorithm is to correct the flow data (segmentation, flow quantification, and background phase error correction). There are 3 flow data-sets that need to be corrected: i) x velocity, ii) y velocity, and iii) z velocity. Due to imaging artifacts (e.g. turbulence) and noise, the flow data will be biased. To correct for this, mass conservation (i.e. physical principles) is used to correct the flow data. Mass conservation tells us that the mass of a closed system must remain constant over time, as system mass cannot change quantity if it is not added or removed. Therefore, if a boundary is defined within the heart (i.e., luminal border of the heart chambers and vessels), the flow entering a stationary volume must match the flow exiting the

volume if the fluid is incompressible. This theory can be applied to any volume. In the case of blood flow, we make assumptions that the blood density is constant, and therefore the continuity equation simplifies to mean that the divergence of velocity field is zero everywhere. Physically, this is equivalent to saying that the local volume dilation rate is zero (i.e., $du/dx+dv/dy+dw/dz=0$). It is impossible to force this condition everywhere, but the local volume dilation can be minimized over all time points. There are several different types of algorithms that will minimize $du/dx+dv/dy+dw/dz$, but the most common is an algorithm that will generate a least squares divergence free approximation of the flow field. There are several ways to construct a least squares approximation to the flow field with the constraint of minimizing the divergence, and several different algorithms to achieve this.

[0100] Typically there is an iterative approach involved that tries to minimize the residual divergence with every pass. In addition, knowing the exact boundary of the vessel/chamber is important to ensure zero flux through the boundary. Without said boundary, flow could be allowed to escape into the heart muscle and fat. In addition, there could be artifact in the image (i.e., caused by turbulence). If a user identifies a region where there is artifact (“bad data”), this region is not used to influence the velocity value correction in the “good data” region.

[0101] Another approach to solve this is using optimization: attempting to minimize divergence while ensuring the original vector field is changed as little as possible (to still capture local effects and prevent smoothing).

[0102] Conservation of momentum can be applied in a later act to estimate pressure gradients across a vessel in addition to wall shear stress. This mass conservation step is critical to ensure accurate pressure estimations.

[0103] Automatic Phase Aliasing Correction Using Time Domain

[0104] Phase aliasing occurs when the VENC that was set for the 4D-Flow scan was too low causing the velocity values to “wrap”; from large positive values to large negative values or vice versa. In principle, this wrapping can happen more than once.

[0105] By analyzing the overall time variation of velocity data it is possible to determine the main points of the cardiac cycle (described elsewhere). Under the assumption that there is no velocity aliasing in the image at peak diastole and also under the assumption that velocity at a single point in space does not, in reality, vary by more than \pm VENC from one time point to the next, one is able to correct for phase aliasing by examining the time variation of each point in space as follows:

[0106] i) Identify the peak diastole time point and assume that there is no phase aliasing at that point.

[0107] ii) Examine the time behavior of each acquired velocity component individually.

[0108] iii) For each point in space (voxel) in each velocity image track the change in velocity from one time point to the next. If the velocity is observed to vary by more than \pm VENC assume that aliasing has occurred.

[0109] iv) When aliasing is detected the wrap count for that point is either incremented if observed velocity reduced by more than VENC or decremented if the observed velocity increased by more than VENC.

[0110] v) At each point in time the velocity is altered according to the current accumulated wrap count for that point by adding the product of wrap count with two times VENC.

[0111] vi) Check that wrap count has returned to a value of zero once the current time point has return to the initial peak diastole starting point. If wrap count has not returned to zero, then the processing of that point in space (voxel) should be considered to be an error.

[0112] The method can be improved and made more performant by using other methods to determine the pixels of interest. For example, one may use other methods to determine the pixels that are most likely to represent blood flow and only process these pixels.

[0113] This method also has the characteristic and advantage of being self-diagnosing. The wrap count for all valid blood voxels (as opposed to air, for example) should return to zero when the processing for that voxel over time has finished. Errors can be kept track of on a voxel by voxel basis though this has the weakness that this method of error detection is not guaranteed to catch every error voxel. However, in addition, by looking for a low overall error rate as a fraction of the number of pixels where corrections were applied, one can ascertain whether or not the necessary initial assumptions, required by the method, are largely correct.

[0114] Automatic Phase Aliasing Correction

[0115] Phase aliasing occurs when the VENC that was set for the 4D-Flow scan was too low. It is very easy to find voxels that have been aliased because of the following:

[0116] i) The VENC of each scan is known since this information is in the header file of all the DICOM images.

[0117] ii) Identify sharp changes in flow velocity around the \pm VENC velocity (i.e., if VENC is set at 100 cm/s, look for sharp changes in velocity around \pm 99 cm/s.) Sharp changes in velocity means a voxel may have a velocity value of 100 cm/s, and the adjacent voxel has a value of -99 cm/s.

[0118] iii) Then find the border of the aliased region by connecting all the voxels that have a sharp gradient around \pm VENC.

[0119] iv) This results in an enclosed boundary. Determine if all the voxels in the region of the enclosed boundary are aliased. By starting at the boundary and moving toward the centroid of the 3D region, the system can interrogate every voxel to ensure that there is no significant jump (across a VENC).

[0120] v) If no steep jump is encountered then the VENC velocity can be added to all voxels within the aliased region.

[0121] vi) If a steep jump is encountered, then the problem becomes a little more challenging (but still solvable). In this case, a voxel could be wrapped several times (i.e., if the VENC is set at 100 cm/s but the velocity at the voxel actually is 499 cm/s, this will be wrapped 2 times and the velocity will be shown as 99 cm/s). The way to correct the data is to look at the velocity of neighboring voxels. If there is a jump of more than 1.5 times the VENC, then $2 \times$ VENC needs to be added or subtracted in that enclosed region. The selection of adding or subtracting is chosen to minimize the discontinuity across the neighboring voxels.

[0122] vii) To improve the algorithm further, information about where static tissue is can be critical in defining where absolute zero must be. Since static tissue has 0 velocity, those voxels that are identified as being static tissue must not be wrapped. There then must be a continuous increase in

speed away from the wall due to physical properties (i.e., fluid boundary layer). The only assumption in all of this is that neighboring voxels do not have more than a 1.5*VENC jump across one another.

[0123] Spline Real-Time Eddy Current Correction

[0124] When an MRI acquisition is performed the data can contain artifacts due to eddy currents in the magnetic field. In a 4D flow acquisition in which particle velocity is acquired the artifacts will cause the velocity values to be incorrect. It is critical with 4D flow and 2D phase contrast ("2D PC") to have accurate velocity data in order to quantify blood flow through vessels.

[0125] At least one technique for correcting the eddy current artifacts involves:

[0126] volumetric segmentation of static (zero velocity) tissues; and

[0127] using the velocity data at these static locations to fit a curve to the volume that can be subtracted from the raw data.

[0128] While it is possible to calculate the model using individual velocity pixels, the resulting matrices are very large and the calculation can take a very long time. To expedite the calculation, the volume may be divided into a grid of blocks, to produce what is effectively a reduced resolution velocity volume. The dimension of the blocks may be adjustable by the user, but is typically 8 mm×8 mm×8 mm, for example. An individual block may then be either masked or assigned a value which represents the average unmasked velocity within the region defined by the block. In order to be assigned a value (in other words, to be considered unmasked) a user adjustable fraction of the region (e.g., 50%) should be unmasked. The model is then calculated using the unmasked velocity values in this reduced resolution volume.

[0129] One method of producing a model of the eddy current error is to fill in the masked regions using spline basis functions wherein the order of the spline may vary due to the number of available control values (with the simplest option being a piecewise linear function between control values). Splines are produced between control points in all three primary directions. The three resulting models are then averaged to produce a single model and the final result smoothed.

[0130] A second method of producing a model of the eddy current error is fitting a low order three dimensional polynomial directly to the unmasked velocity values in the reduced resolution volume.

[0131] Because the resolution of the volume has been greatly reduced, both of these methods are efficient options and can be easily executed at run time.

[0132] Once a reduced resolution model of the eddy current error is produced, it is up-sampled to the original resolution and then subtracted from the original data. After subtraction, the static tissues should have an effective velocity of zero, and the non-static tissues will have eddy current artifacts removed. This will allow for accurate flow quantification. Based on the analysis of our test data, results from this new and innovative approach is a smooth 3D time-varying function representing the error in the volume and is very fast to calculate.

[0133] Based on the analysis of our test data, results from this new and innovative approach is a smooth 3D time-varying function representing the error in the volume and is very fast to calculate.

[0134] For the final step of upsampling to the original resolution, our approach to applying the correction algorithm was to use tri-linear sampling from the correction volume, which proved successful. For each element in the source volume we perform a tri-linear blend of eight elements in the correction volume and subtract that value from the source data. After applying our new correction function flow measurements were found to be within 3% error. In addition, as part of the process requires user interaction, the requirement for real-time performance was also met as evaluating and applying our new correction takes on the order of milliseconds instead of hours.

[0135] Solids for Background Phase Error Correction

[0136] The blood velocity information in the flow-4D MRI scan has an error which needs to be corrected for in order to get accurate blood flow calculations. The error signal in the static tissue can be used to provide a correction function for the non-static tissue (described elsewhere). A three-dimensional volume of tissue called a solid can be created to identify a section of either static or non-static tissue. A solid can be created using two methods:

[0137] Orthogonal contours: The user can manually draw some number (typically three and typically orthogonal) intersecting cylinders defined by closed contours. The intersection of these contours represents a solid, three-dimensional volume of either static or non-static tissue. The contours do not need to be perfectly orthogonal and the user can create the contours at any location.

[0138] 3D floods: The user can alternatively choose to automatically create a solid by specifying the starting point of a three-dimensional flood. The flood can be used on any image including phase images. The image is flooded based on a threshold below and above the value at the point where the user clicked. The user can control both the threshold and the maximum radius of the flood that is generated.

[0139] Multiple solids can be created using either method to mark areas of both static and non-static tissue. The masks associated with non-static tissue are added together. The masks associated with static tissue are then used override or unmask previously masked regions as necessary.

[0140] In some situations there will be artifact in the images. Artifact needs to be removed or highlighted in order to prevent a user from giving an incorrect diagnosis. Our software has pre-processing steps (i.e., eddy current correction) that compute a metric based on all of the voxels in the dataset. To avoid these pre-processing steps a tool is used to identify voxels that have artifact, and these voxels are removed from any pre-processing step (but not removed for visualization purposes). The tool can be a manual segmentation tool (i.e., user circles areas that have artifact), or a semi-automatic/automatic segmentation tool that identifies artifact. Regardless of the tool specifics our software needs a feature to remove "bad voxels" from being quantified.

[0141] Automatic Background Phase Error Correction

[0142] Accurate measurement of velocity in flow-4D MRI scans requires the application of corrections for the false signal introduced by eddy currents. Determining eddy current correction (ECC) is done by examining the velocity signal in static (non-moving) tissue. This requires the masking of all moving tissue, blood and air. This claim describes a method for doing this automatically, without user intervention. Automatic correction is useful as it not only makes the software simpler and quicker to use, it allows the correction to be precomputed and applied when the user first

opens the study. It is also very important as it allows other preprocessing algorithms, doing things like automatic segmentation and measurement, to benefit from ECC.

[0143] Automatic ECC is done by computing initial starting values for three filters that the user is then free to adjust after the study is open. Air is masked by masking out regions with anatomy image values below a set threshold. This threshold is determined automatically by an analysis of the histogram for anatomy image values over the entire scanned volume. The histogram for these values over the chest cavity displays a pattern that allows for the automatic detection of the image values corresponding to air.

[0144] In addition, two filters have been developed that reliably detect regions of blood flow and regions of heart wall movement. The region of the heart can be satisfactorily masked out by appropriately setting these two filters. Due to the normalization used in the production of these filters along with their naturally consistent nature, satisfactory results can be obtained simply by setting these filters to predetermined values (for example, 50%). The automatic setting of these filters could be further improved (or tweaked) by analysis of the values produced by these filters, similar to what was described in the previous paragraph for the detection of regions of air. The settings could also be tweaked by examining the resulting ECC and looking for regions that show large variation over the cardiac cycle.

[0145] The correct values for these filters, determined when the study is preprocessed, are then simply stored in the database along with the other study information, allowing the client software to set these as the default values when the study is first opened.

Visualization

[0146] Temporal Landmark Quantification and Visualization

[0147] It useful to be able to identify landmarks (i.e., points, lines, planes, areas, volumes) in the body and specifically in the heart. Several landmarks are dynamic in nature (e.g., mitral valve plane), and thus it is important to track their movement in time:

[0148] Points: Track 3D path (which is a line) over time.

[0149] Lines: Track 3D path of 2 endpoints over time.

[0150] Planes: Track a point on a plane and the plane's normal vector over time.

[0151] Areas: Track dilating contour, contour centroid, and contour normal vector over time.

[0152] Volumes: Discretize surface of volume and track each discretized point over time.

[0153] There are two acts to temporal landmark quantification and visualization:

[0154] 1) Detection: The first step is to identify the landmarks over time. This can be done manually or automatically.

[0155] Manual detection: The user can indicate the position and orientation of each landmark. One method of doing this could be to navigate the image using pan and rotate so that the center of the image is at the desired location of the landmark. The location of the landmark can be different at different points in time and it will be interpolated for time points where the user has not explicitly set it. It is indicated to the user if a landmark is interpolated.

[0156] 2) Display: Depending on the type of landmark a different type of method to display the data is used. For instance if a contour does not move or change it's normal

vector over time (i.e., just dilate), it makes sense for the user's view plane to not change and always be aligned with the contour. If this contour does move, we can imagine following the plane such that the view is always aligned with the plane for each time point. The view plane can be either from a Lagrangian perspective or from an Eulerian perspective. For volumes, an Eulerian perspective is more appropriate where the surface of a volume dilates and this can be visualized with a camera that is fixed in space (the user can change the camera location as needed).

[0157] Cardiac views: The left ventricle apex, right ventricle apex, mitral valve, tricuspid valve, aortic valve, and pulmonic valve can be used to create two-chamber, three-chamber, four-chamber, and short-axis views of both the right and left ventricles once the landmarks have been detected. The orientations of these views are specified in the Mayo Clinic Guide to Cardiac MR. An orientation and zoom level for each view can be calculated from the positions of the landmarks. If the landmark's position changes in time the view will change in time accordingly.

[0158] Example landmarks for each view:

[0159] Left two chamber: aortic valve, mitral valve, tricuspid valve, left ventricle apex

[0160] Left three chamber: aortic valve, mitral valve, left ventricle apex

[0161] Left four chamber: tricuspid valve, mitral valve, left ventricle apex

[0162] Left short axis: mitral valve, left ventricle apex

[0163] Right two chamber: pulmonic valve, tricuspid valve, mitral valve, right ventricle apex

[0164] Right three chamber: pulmonic valve, tricuspid valve, right ventricle apex

[0165] Right four chamber: tricuspid valve, mitral valve, right ventricle apex

[0166] Right short axis: tricuspid valve, right ventricle apex

[0167] Interactive Landmark Based Views

[0168] Once certain landmarks have been placed (e.g., aortic valve, mitral valve, left ventricle apex, anterior papillary muscle, posterior papillary muscle, pulmonary valve, tricuspid valve, right ventricle apex, LPA, RPA, SVC, IVC, descending aorta), automatic views can be created to display the anatomy of interest. Clinicians are used to viewing a certain landmark with 3 perpendicular views or in the case of the heart using a 4 chamber view or a left or right ventricle 2 or 3 chamber view. By updating the location of just 1 landmark for one of the time points, all views update accordingly such that either the views are always perpendicular or the 2, 3, and 4 chamber views remain intact. Once the landmarks have been placed, and the views generated automatically, these views can be saved in the report section of the software and exported in any format (i.e., an image) include a cine movie (i.e., multiple images over time).

[0169] 4D Mesh Creation from Closed Contours

[0170] Once contours have been placed along the short axis (possibly curved) for each time point, the mesh is then generated independently for each time point. This is done by rotating each contour in the short axis stack so as to minimize twisting, and then generating open cubic spline which connects the first point in each contour, a second spline that connects the second point, so on for each point in the contour (each slice's contour has the same number of points). The result of this process is a cylindrical grid of points which we use as the vertices of the mesh.

[0171] The process of minimizing twisting is done by computing an open cubic hermite spline from the centroid of one contour to the centroid of the contour above, and then running this spline from each point on the lower contour until it intersects the plane the contour above it lies in. The system computes this intersection point and then determines which of these intersection points lies closest to an actual contour point in the upper contour. The contour is then rotated such that these two points will lie on the same long axis spline.

[0172] The current implementation works reasonably well with both curved and straight axes when the contours are reasonably circular and the difference between neighboring contours spatially is minimal. However, in the case of the right ventricle where the contours are not circular, excessive twisting is sometimes introduced by the current implementation. To minimize this, we should get rid of the long axis spline approach and switch to one where the number of triangles between any two slices may differ. Doing this minimizes twisting more locally which will result in an overall smoother mesh.

[0173] Snap to Flow Tool

[0174] Accurately observing or measuring the blood flow in a flow-4D scan requires the user to align an MPR so that it is perpendicular to the direction of flow. This describes a method for creating a tool allowing the user to set the correct orientation of the MPR automatically.

[0175] In order to align an MPR the user first activates the tool and then clicks on a central region of the blood flow in question. The click point then serves as the center of rotation when the MPR is aligned, moving the click point to the center of the resulting MPR. Alignment is done by averaging the blood flow in a small region around the click point. To do this accurately, the measurement is done using the time point corresponding to peak blood flow, regardless of the time point that the user is currently viewing while using the tool. This generally implies doing the measurement at peak systole.

[0176] While the user is allowed to adjust the time point for peak systole, this point is first determined automatically by the executing software during preprocessing of the data set, and this automatic value is used as the default when the study is first opened by the user. A filter has been developed (described elsewhere) for automatically determining regions of blood flow within the scanned volume. Peak systole is then determined by examining the time dependence of the overall flow within the filtered or mask region determined to correspond to blood.

[0177] Once the direction of flow has been accurately determined it is straightforward to adjust the orientation of the MPR so that it is on a plane perpendicular to the flow.

Quantification

[0178] Automatic Blood Flow Quantification

[0179] The blood flow in a chamber and/or vessel can be automatically quantified by first isolating the blood pool (see segmentation methods described in this document) and placing a plane on a landmark (that can be defined using the methods above) that is roughly perpendicular to the flow in the chamber/vessel (i.e., normal of plane is aligned with the flow). Once these 2 acts have been achieved, the intersection between the plane and blood pool creates a contour. All the voxels within the contour are flagged. Next is to sum the dot product of the plane's normal vector with the velocity vector

of that voxel (in addition to normalizing by the area of the voxel) for every voxel to give the total flow. The flow at that contour can be automatically displayed on screen or in report that could be eventually exported.

[0180] Allowing a user to select a position on an image has many important applications. In performing measurements a user might want to measure the distance from one point to another. In an application that uses MPRs from a volume of data, the points on an image represent locations in 3D space. These 3D points are easy to compute from the metadata associated with the image. In an application that uses volume rendering, allowing a user to select a point in 3D space is more difficult since each pixel could be at a different depth.

[0181] In typical front to back volume raycasting with an increasing alpha compositing function that terminates once alpha reaches 1.0, determining the depth of the pixel can be done by keeping track of where the ray terminates. When raycasting back to front, there is no early ray termination. The result color is simply updated based on the compositing function. Typically the compositing function will make air transparent and as such the color will stop changing as the ray exits the material closest to the eye. By keeping track of when the color stopped changing, this depth for each pixel can be used to transform the 2D user selected coordinate back into a 3D location in space. This 3D location selection can be used to select a blood vessel and then automatically quantify flow.

[0182] Automatic Shunt Detection

[0183] Instead of trying to find an exact location for a shunt, the first operation would be to identify if a shunt exists. One simple method of identifying if a shunt is present is to measure the left heart flow (Q_s) and the right heart flow (Q_p). Q_p and Q_s can be measured either manually (e.g., by placing a contour) or automatically if landmarks and blood pool segmentation have been completed. If these numbers do not match within a certain threshold, the scan can be flagged as potentially having a shunt.

[0184] These measurements could be done automatically using the following technique:

[0185] i) Automatic measurement of cardiac output (Q_s) is described elsewhere as is the production of masks for both aortic and pulmonic flow along with automatic estimates of the location of the aortic and pulmonic valves.

[0186] ii) Once the valve regions have been identified, it is a straightforward task to take them and the already determined pulmonic flow regions, move slightly downstream from the valve and produce flow measurement contours in a similar way to what has been described for cardiac output. Once suitable contours have been identified for measuring pulmonic flow the existing flow measurement algorithms can be used to determine the output from the right ventricle.

[0187] iii) Use the automatic flow measurement to indicate the likelihood that a shunt exists.

[0188] Automatic Detection of Peak and End Systole and Diastole

[0189] Much automatic processing depends on the ability to first identify the time points corresponding to the main temporal landmarks in the cardiac cycle: peak and end systole and diastole.

[0190] As described elsewhere, we are able to use a Fourier analysis technique on the velocity images in order to identify regions of blood flow within the heart along with the

main arteries and veins around the heart. Once these main regions of blood flow have been identified, we find the total blood flow over the identified voxels at each point in time (typically 20 time points). The system is then able to analyze the resulting function of time to determine the landmarks in the cardiac cycle. The time point with the most flow is first assigned the peak systole landmark. From there the function is analyzed in both directions in time to determine the points where the flow tends to level off. The point before peak systole where the total flow levels off (the point right before it starts to rise quickly) corresponds to end diastole. Following peak systole, the total flow drops quickly until it levels off, which corresponds to end systole. Peak diastole is not typically a well-defined point so we place this temporal landmark at the point midway between end systole and end diastole.

[0191] Automatic Cardiac Output and Volumetric Measurements

[0192] Automatic measurement of cardiac output is done using the following method:

[0193] i) The relationship between the main DFT components of the velocity images, along with the already determined peak systole landmark (described elsewhere) are used to identify the main regions of arterial flow from the left and right ventricles.

[0194] ii) A variety of flow continuity filters is used, one after the other, to separate the arterial flow region into two pieces, aortic and pulmonic flow. The point in the initial arterial flow mask with the largest velocity provides a reliable point known to be in either the aorta or the pulmonary artery. Separation of the two regions of flow can be determined, for example, by examination of the size of the region within the resulting filter that can be flooded starting at the point of maximum flow. Once the first piece is identified, the second piece can be identified, for example, by flooding from the maximum flow point in the remaining regions.

[0195] iii) Once two regions have been identified one corresponding to aortic flow and one to pulmonic flow, the two regions can be allowed to grow back a limited amount (with individual pixels only being assigned to one mask or the other) and with the original arterial flow mask providing absolute limits to the amount of growth. Allowing at least a little dilation of the masks can also be very important as the preceding process operations may have put small holes in the resulting regions that would tend to hinder the next steps in the method.

[0196] iv) The two flow regions can be identified as aortic and pulmonic flow based on their spatial relationship to each other and their very different expected shape and orientation in space. Once this is done the original arterial flow mask is essentially divided into two regions, one labeled aortic flow and the other labeled pulmonic flow.

[0197] v) As the aorta is essentially one continuous pipe, the path of the aorta can be traced from a starting point within the artery until the two ends are reached. At each point the main peak systole flow direction can be determined by averaging over a small region around the point. Orthogonals to the flow direction can then be projected from the starting point at regular angular intervals to determine the boundary with the masked aortic region, thus determining an approximately circular contour around the starting point.

[0198] vi) Once a contour has been determined as a polygon on the plane orthogonal to the main flow direction

for some starting point. The starting point is re-centered in the polygon. At this point a small step (for example, one millimeter) can be taken from the center point in the positive or negative flow direction, depending on which way from the starting point we are tracing, and the process then repeated. This is continued until we step out of the mask at each end.

[0199] vii) Once contours have been produced at regular intervals along the aorta, essentially producing a mesh, they are refined at each individual time point using either the anatomy images (possible if dealing with a blood flow enhanced dataset) or by using through velocity for the systole time points and interpolation between. One possible approach is to use a snake algorithm to accurately identify the desired boundary for each contour at each point in time.

[0200] viii) Once refined contours have been determined, the major and minor diameters of each contour are measured, the major diameter being the largest diameter and the minor diameter being the largest diameter orthogonal to the major diameter.

[0201] viii) The next task is to identify good contours in the main region of the ascending aorta between the aortic valve and the bifurcations that occur at the top of the aorta, as this is the region that needs to be used when measuring cardiac output. This can be done in a number of acts. First, ascending aorta regions are easily separated from descending regions by flow direction. The remaining contours can then be scored using a combination of the continuity and variability of the contour area and diameters (major and minor) both spatially (along the aorta) and temporally at one point in the aorta. Scores can be averaged along the aorta to look for regions of good scoring as opposed to simply identifying individual, highly scored, contours. Using this method, one can eliminate regions in the neighborhood of the bifurcations at the top of the aorta and also regions that might exist near the aortic valve and on into the left ventricle, as these regions, by their nature, will score badly.

[0202] ix) Once good regions of the ascending aorta have been identified, the highest scoring individual contours can be selected for the actual cardiac output measurement. If possible, measurement is done at multiple points along the ascending aorta, which improves the result through averaging along with providing automatic determination of the quality of the measurement by examining the variability (thereby, also providing estimates of the measurement uncertainty). In addition, examining the result of multiple measurements of the flow along the ascending aorta allows for a judgement on the quality of the velocity eddy-current-correction that is currently being applied.

[0203] x) Once the ideal contours have been selected along the ascending aorta, cardiac output is determined by the usual flow measurement techniques.

[0204] 4D Volumetric Measurement

[0205] In order to calculate the volume of a particular region, we have developed four options within the Arterys cloud interface.

[0206] Option 1: Fixed Axis

[0207] Two points in 3D space define the primary axis of a volume of interest. A straight line connects these 2 points (i.e. fixed axis). The axis is then divided into discrete points (say 2-40 for example) that define the locations where a slice will be placed. Slices are aligned orthogonal to the axis such that they do not intersect. Slices do not have to be evenly spaced. An MPR is rendered at all slice locations to

allow a user to see what the medical image looks like at that slice location. Then either manually or automatically a closed contour is created on every slice to define the boundary of the volume at that slice location. There could be multiple closed contours at every slice location. There could also be no contours at all on one or more slices. In the case of 4D or higher dimensional studies (i.e., studies that show volume change, or said different, multiple frames per slice), there can be separate contours per frame. Once all contours have been placed for all frames and slices, a 3D surface is created connecting the contours of all slices for a particular frame. The way the 3D surface is created from a set of closed contours is explained above in “4D Mesh Creation from Closed Contours”. If there is a 4D or higher dimensional volume, a change in volume can be calculated by computing the volume of each frame and subtracting it with another frame. This is especially important when trying to quantify ventricular function which then gives stroke volumes and ejection fractions.

[0208] Option 2: Moving Straight Axis

[0209] This method is similar to Option 1, except that in the case of a 4D volume, the landmarks or points that define the two endpoints of the axis can move over each frame (e.g. time point). This causes the volume to potentially move locations in 3D space without changing volume.

[0210] Option 3: Fixed Curved Axis.

[0211] This method is similar to Option 1, except that the line connecting the 2 endpoints does not have to be straight. This line can be curved or have multiple straight and curved sections. This is handled in the system with a spline that connects points/locations between 2 endpoints. These points/locations can be anywhere and not necessarily always between the 2 endpoints.

[0212] Option 4: Moving Curved Axis

[0213] This method is similar to Option 2, except that in the case of a 4D volume, the landmarks or points that define the two endpoints of the curved axis can move over each frame (e.g. time point). This causes the volume to potentially move locations in 3D space without changing volume.

[0214] In all of the options above, there could be multiple axes. For example there could be a “Y” shaped axis that splits in 2 from 1. There is also the option of having both straight and curved axes that split and come together to create the volume. The point of this would be to account for more complex shapes that still have a primary axis (i.e., centerline).

[0215] In all of the options above, there is also the option of displaying how the 3D volume intersects an MPR. The intersection must be a collection of one or more closed contours. These closed contours can be rendered on the MPR. In addition, these closed contours can be edited by moving the contour in the new (non orthogonal) view. The intersection contours can be computed both on the client as well as the server, or be adaptive depending on local resources. For cardiac imaging, common non-orthogonal views are 2, 3, and 4 chamber views. The contours can be edited in these views by only allowing the editing to be in a certain direction (i.e., along the slice plane).

[0216] Out of Plane Measurements and Tracking Mode

[0217] Measurements made on structures in a cardiac system from volumetric MRI data, such as flow measurements of a heart valve, must deal with several complexities such as the shape, position, orientation, and velocity of the structure changing significantly over a cardiac cycle. For

measurements, such as flow, where the measurement is determined by a 2D contour, we solve this by using 2D contours that move through 3D space. In the case of a heart valve, either manually or automatically, contours are placed at the border of the valve opening on the plane that is most perpendicular to flow direction. The position and orientation of the valve plane are tracked for each phase of the cardiac cycle with the measurement contour placed accordingly. The evaluation of flow may be performed through standard finite methods integration. However, in the event that the valve plane is moving, the linear and angular velocity of the valve plane can be included in the flow computation for that phase. During visualization, when cycling through phases, the position and orientation of the displayed MPR can track the valve plane following the position and orientation of the valve as it moves. If a measurement is visualized when the current MPR is out of plane, the contour may be rendered semi-transparent.

Segmentation

[0218] Continuity Equation Driven Segmentation of Blood Pool

[0219] Once again, mass conservation (i.e., continuity) with the incompressibility assumption can be used to show that divergence must be zero everywhere in the blood pool. By computing the divergence everywhere, the system can define the extents of the blood pool by a threshold divergence value. The divergence outside the blood pool will be larger (i.e., air in the lungs) or the velocity will be low (i.e., velocity signal in static tissue), which both help in identifying the lumen boundary. The divergence map does not need to be the sole input into a segmentation algorithm, instead it could be added to other inputs and weighted appropriately.

[0220] Automatic Landmark Detection

[0221] The typical ways to create an automatic landmark detection algorithm is look for certain shapes in images and measure distances and angles between these shapes. If the measurements lie within a certain band they are classified. Several other physiologic inputs can be added to the algorithm. For instance locating a volume of fluid that increases and decreases substantially with every heartbeat (this is likely to be a ventricle). Once a ventricle is found, the inlet and outlet of the valve can be found by following streamlines. Once a valve is found, it is easier to find the remaining valves because they are typically always a certain distance and angle away from each other.

[0222] The algorithm that is selected to find the landmarks can be of the machine learning type. Since Arterys will be constantly collecting data that has been validated with correct landmark placing by a clinician this data needs to be used as a training set (e.g., statistical aggregation of data). Every dataset that needs to be analyzed can then be co-registered with an ‘atlas’ that is built with the training set data. Once a sufficient number of datasets are collected, additional input parameters such as type of disease (i.e., healthy, tetralogy of fallot, etc.) can be used to bin the datasets prior to be analyzed. Every bin could have slightly different landmarks and measurements depending on the type of disease and what pathology is expected. If it is known that a dataset is a patient that has a single ventricle, the automatic landmark detection algorithm needs to adjust for this as it will never find 4 valves.

[0223] In particular the aortic and pulmonic valve landmarks can be determined using the following process:

[0224] i) Identify regions corresponding to arterial flow from the left and right ventricle. Filters have been developed (described elsewhere) that are able to do this with high reliability.

[0225] ii) Separate the region of arterial flow into two regions, one corresponding to the aorta and one to the pulmonary artery. This process is described in detail under cardiac output.

[0226] iii) Once one region corresponding to either flow from the left or right ventricle has been determined, the other region is determined by subtracting from the starting region corresponding to both flows. The regions can then be easily identified as left ventricle flow or right ventricle flow based on their physical dimensions and orientations in space (also described under cardiac output).

[0227] iv) Once the two regions of flow have been identified, initial approximations for the location of the aortic and pulmonic valve can be determined by carefully tracing the bulk flow back to its apparent origin.

[0228] v) Once reliable initial estimates are produced for the location of the two valves, other techniques can be used to refine the valve location. For example, one could examine the blood flow acceleration and intensity in the region surrounding the initial estimate in order to refine the location of the valves.

[0229] Interactive 4D Volume Segmentation

[0230] Segmentation of ventricles from a cardiac scan is critical to determining ventricular function. Automatic ventricular function technique may involve:

[0231] an input of two or more points representing control points of a spline;

[0232] the endpoints of the spline denote the apex of the ventricle and the exit valve (pulmonic or aortic);

[0233] using these points generate MPRs with plane normals set to the tangent of the curve at regular intervals along the spline curve;

[0234] on each MPR apply an active contour model to find the boundary (epicardium or endocardium) of the ventricle; and

[0235] generate a 3D mesh using the points of each of these contours.

[0236] Active contour models are subject to instability from the forces that act on them. To reduce this instability, instead of simply generating the contours such that they are spaced at the desired output spacing (distance between contours), the system generates many contours spaced very tightly together. Also, if the input data has temporal data, contours at the same location are generated using data from adjacent time points. Contour shape and quality is then measured against typical contours from a ventricle. If a contour is deemed to be of sufficient quality it is included in generating a final result. The final results are generated by averaging the included contours that are close to the position and time along the input curve. With a mesh constructed at both end systole and end diastole the difference in volume represents cardiac output and ventricular function.

[0237] In one example implementation, the Arterys system and software would provide single click 4D volume segmentation. This would allow the user to click areas of interest (e.g., blood pool, myocardium, bone, etc.) while freely navigating (i.e., rotating, panning, zooming, slice scrolling, time scrolling) the 3D volume. Since a full 3D

volume segmentation algorithm is challenging to construct and be accurate, a second option is to display 3 orthogonal views to the user while the user draws the boundary of the area the user would like to segment. For the heart, the view that is displayed can be a 2, 3, and 4 chamber view of the heart in addition to a short axis view. The user only needs to create 2 orthogonal contours in long axis, and then the software can automatically or autonomously create a 3D surface based on interpolating the two contours. The 3D surface can be shown in short axis to the user for quick modification. In addition to showing the anatomic images, the blood velocity images (with or without vectors) can be overlaid onto the anatomic images to further clarify where the blood pool boundary is during the interactive 3D volume segmentation process.

[0238] Adaptive Flood Fill

[0239] The system makes use of multiple types of floods which may be distinguished as 2D vs. 3D, by connectivity used during the flood (6, 18, or 26 way connectivity), and radius constrained vs. a flood constrained by a maximum number of steps. In all cases, the flood works by moving outward from a specified seed point and including a pixel in the result of the flood if it is 1) connected to the rest of the flood (using whatever connectivity was specified), 2) has an intensity within a specified threshold of the pixel at the seed point, and 3) the pixel is within the specified radius or maximum number of steps of the seed point. The result of the flood is a 2 or 3 dimensional connected mask. The flood algorithm is used in solids in the form of a 3D flood to mark static/non-static tissue, in volumes where a 2D flood can be used to generate a contour in the short axis stack, and in flow quantification, where a 2D flood may be used to flood a vessel to determine the flow contained within the flood.

[0240] To generate a contour from a 2D flood, we make use of the fact that the flood will necessarily be connected and that it is a binary image. Because of these facts, we may apply a standard border tracing algorithm to come up with a contour which will ignore any holes that may be present within the interior of the flood.

[0241] From the generated contour, the next operation is to reduce the generated contour from potentially hundreds of points to a small set of control points to be used by a closed cubic spline to accurately approximate the actual contour. A naïve down sample where the system simply spaces a fixed number of control points spaced equally around the contour does not work as well as other approaches, as this approach frequently results in the loss of important features in the contour such as concave portion of the flood which was going around a papillary muscle. To get around this, a “smart” down sample approach is employed which proceeds in a number of acts. To begin with, each point in the contour is assigned a corner strength score ranging from -1 to 1, as well as assigning each point an area of “influence”. Once this is done, the contour is reduced to only those points where their corner strength is maximal within their area of influence. Additional criteria are also enforced in this stage, such as ensuring we have a minimal point spacing and ensuring our detected corners are sufficiently strong.

[0242] The result of the preceding operation is a list of “corners” detected in the flood. By using these as control points in a spline, this approach ensures that the spline will not lose any interesting features in the contour. However, any long stretches of relatively low curvature in the original contour will not be detected as corners, which can result in

significant portions of the resulting contour not having any control points, leading to a poor approximation by a spline in such segments. To get around this, an error metric is computed for each pair of control points by calculating the area of a closed contour formed by the segment of the original contour passing through the points, and the segment of a spline passing through those points. If the error is above some fixed tolerance, another control point is added at the midpoint of the segment of the original contour. This operation is repeated until each segment has a computed error below the required tolerance.

[0243] This flood-to-contour tool can be used in at least two places in the application: for flooding slices of a ventricle while performing a volumetric segmentation, and in flow quantification. In the case of the flood for volumes, the returned contour is dilated by 8% in order to capture more of the ventricle as a raw flood fill often underestimates simply because of the difference in pixel intensities close to the heart wall. For a flow flood, the result is dilated by 12% because the flood tool works on anatomy, which means the undilated flood will often miss flow near the vessel wall.

Overall Process

[0244] Automated Reports

[0245] In a way similar to how echocardiographic reports are generated, an automated report based on 4D-Flow MR data can be created by allowing the user to click on the type of patient they have. Arterys will have unique report templates that are specific to a certain pathology or type of user (i.e. patient or clinician). All of the values, curves, images, and cine movies in this report can be automatically populated in the report template. Since landmarks are placed as part of the pre-processing step, all the important information can be automatically saved in the database and exported to this report.

[0246] Automated Integration Tests

[0247] A tool called nwjs that is designed for producing cross platform native applications using Web technologies is used to perform automated integration tests. Although not designed for this purpose, it allows us to run both client and server software stack within the same environment allowing us complete control over the client and server applications at the same time. Using infrastructure mixed with a test tool called mocha, we write tests that emulate the customer interaction with the client while asserting both the client and server processing of that interaction along with the resulting state of the application. This method of integration testing is novel and superior to other tools that are mostly vision based, for this type of User Interface testing.

[0248] Hybrid Client Server Rendering

[0249] Description Some workflows require one or multiple images to be rendered at the same time that have linked properties. In some cases the current workflow step may require simultaneous viewing of 20 images. If each of these images was retrieved with a distinct HTTPS request, performance would suffer greatly as there is significant overhead in creating and sending a request. Instead, we render all the images onto one large image, and only make a single HTTPS request for that 'sprite sheet'. The client then displays the images by using pixel offsets. For example, if a view had four images each 256×256, the sprite sheet might be 256×1024 with each of the images stacked one on top of another. The client would then display 4 images at 256×256 by using offsets of 0, 256, 512, and 768.

[0250] In addition, any lines, markers, or planes in the images are drawn on the client as an overlay, and the information that informs the client how to render the overlay comes from the server via a JSON message. This provides higher quality rendering of the overlay data than if the overlay were to be rendered on the server and then encoded as a JPEG and transmitted.

[0251] Automated Global Endurance and Stress Testing

[0252] In order to perform load testing and endurance testing, we launch a multitude of client processes on a multitude of computers (which can be geographically distributed) to start specialized web browsers in which we have complete control over their execution environment. They are directed to the application and load the client as a normal browser would, then we directly interact with the client state controlling the software and making it behave as certain workload. The client and server metrics are recorded during load testing and can be run for longer periods of time for endurance testing.

[0253] Pusher Pushing Data from a Medical Device to Remote Servers

[0254] We have developed software to monitor for active studies, and push the results to our remote service in the cloud. A folder is monitored for files being generated by a scanner, and upon completion, all relevant data is bundled together and pushed via a secure connection using a unique secret and key per scanner for authorization to our remote cloud servers. Disk space (e.g., nontransitory storage media) usage is minimized by immediately deleting any intermediate files.

[0255] Upon a successful transfer, the data integrity of the transferred content is verified against the local content by reproducing the package process and comparing the output of a cryptographic hash function. Repeating the process like this ensures that any new data that may have been generated by a scan was not missed in the case of delays during the scanning process which may trigger premature transfer of the data to Arterys servers.

[0256] In the case of a failed transfer, due to server or network errors a configurable number of attempts will be made with an increasing interval of rest between each attempt, before the pusher assumes the transfer was a failure. However, after a failed transfer (including all subsequent attempts), the pusher will continue to monitor for incoming files, and will re-attempt another transfer at a later time.

[0257] Once data has been verified as successfully transferred, the data is removed by our software to conserve disk space on scanners.

[0258] A heartbeat message is sent from each pusher software running on every scanner providing the local log data and detailed status information of the scanner, providing continuous monitoring and increased response time to ensure scanner functionality during critical scan times.

[0259] During initial installation, a scanner will automatically register with Arterys by requesting a unique secret and key to sign all future requests with for authorization purposes. The scanner will be registered in our systems database, but not attached to any organizations. A technician is then able to attach all recently registered scanners to the correct organization through a web portal.

[0260] A pusher is able to auto update (if configured) by periodically requesting new versions from Arterys. If a new version is provided, it will install a new copy of itself, and restart. This allows for security and functionality updates to

be deployed to scanners with no intervention from technicians. The heartbeat messages provide the information required to ensure success of this operation on the Arterys servers. The heartbeats enable us to determine any pushers that have not been updated recently, and reach out to hospitals directly to proactively ensure all software is up to date and secure.

[0261] FIG. 3 shows an example process 300.

[0262] Puller—Archiving Artifacts

[0263] The puller software is used to archive generated artifacts at a hospital (for example PACS). It is installed within a hospital's network and registers itself automatically with Arterys using a similar method to the pushers. A request is made with some identifying information, and a secret and key pair is returned to sign future requests for authentication and authorization purposes. The puller is then attached to an organization by a technician through a web portal.

[0264] It is also possible to download a version for an organization directly, with a unique key and secret included automatically in the installation process, so there is no need to auto register and attach the puller once installed.

[0265] The configuration for artifact endpoints is done on Arterys servers. Multiple locations can be configured with hostnames, ports, AE titles, and any other required information the puller would need to transfer data to it. These endpoints can be named, and are selectable from the Arterys web UI by a clinician when choosing where they would like their artifacts (reports/screenshots/videos) to be archived.

[0266] The puller monitors for artifacts by requesting a list from the Arterys API at regular and frequent intervals. The list of artifacts includes a unique id, and all of the configuration information for the endpoint the artifact will be stored in. The unique ID is used as input into another API request to retrieve the artifact from Arterys servers. The artifact is unzipped if required, and transferred using the configuration and method defined by the configuration included in the list request (e.g., store scp). Once all data is transferred, another API request using the provided ID is made to Arterys to mark the artifact as archived, and it will no longer appear in the list generated by the first request in the process loop. Once the artifact has been marked as archived, the Arterys servers will notify a user that archival is complete.

[0267] The puller sends heartbeat requests to Arterys providing detailed logs to help validate and ensure everything is functioning as expected. The puller will also occasionally—at a configurable time (e.g., once an hour or day)—make an API request to Arterys for new versions of the puller software. If a new version is available, it will be downloaded, installed and the puller will restart itself.

[0268] Example request to retrieve a list of artifacts

```
GET https://app.arterys.com/api/1/artifact?status=pending&timeout=20
```

```
[
  { "id": "55006362619baaad4323f799", "name":
    "report_one.zip", "digest": "f90sdaF9d0safd0safd09safd",
    "size": 3523234, "dicom_host": "192.168.1.3",
    "dicom_port": 1999, "dicom_aetitle": "aetitle for dicom
    endpoint" },
  { "id": "454bf977be1cfbe146f36549", "name":
    "report_two.zip", "digest": "9320028003002930sass9safd",
    "size": 1221134, "dicom_host": "192.168.1.3",
```

-continued

```
GET https://app.arterys.com/api/1/artifact?status=pending&timeout=20
```

```
"dicom_port": 1999, "dicom_aetitle": "aetitle for dicom
endpoint" }
]
```

[0269] FIG. 4 shows an example process 400 of monitoring for artifacts and archiving.

[0270] PHI Service

[0271] We have developed a method to securely deliver sensitive patient information to a client application from a service without disclosing the sensitive information to the service provider.

[0272] The data prior to being sent to the service provider is stripped of all patient identifiable health information, which is registered with the service and the original sensitive data is replaced with a unique token identifier provided by the service.

[0273] The client when interacting with the service provider will identify these tokens and use an independent transport layer to replace the tokens with the sensitive patient health information.

[0274] Below is an example of a possible implementation of such a system:

[0275] Actors:

[0276] The user which integrates with the client software (user)

[0277] The client application (client)

[0278] The service which holds the sensitive patient information (service)

[0279] The application service provider (asp)

[0280] 1. The user indicates to the software a set of files it would like to send to an asp.

[0281] 2. For each file sensitive all sensitive information is gathered in json format and registered with the service over an http request.

Example

[0282]

```
POST https://sensitive.arterys.com/register-data HTTP/1.0
{
  "PatientName": "Franklin\Benjamin"
  "Birthdate": "1706-01-17"
}
returns
  Location: "/4217ad2b78ff7eb9129a58b474efb3e"
```

[0283] 3. The sensitive data is replaced with placeholders such as `#{PatientName}` and then the data is uploaded along with the Location url returned from the service.

[0284] 4. When the client loads the data from the application service provider, strings that contain these sensitive tokens cause the client application to request the data from the service provider (either individually or in bulk).

Example

[0285] GET
<https://sensitive.arterys.com/4217ad2b78ff7eb9129a58b474efb3e#PatientName> returns
 [0286] “Franklin/Benjamin”
 [0287] 5. The client substitutes the tokens with the sensitive information.
 [0288] Note: For authorization, we could use an sso such as saml2.
 [0289] Workspaces
 [0290] Workspaces are a solution the issues of storing and sharing a subset of application state throughout medical software.
 [0291] Workspaces contain the application state of a study including any analysis, and when loaded they restore application the previous state. Application state includes the subset of component state related to a particular concern such as study review including measurements and ecc correction values etc.
 [0292] Workspaces can be loaded and updated constantly while the user interacts with the software. Users start with a private default workspace when loading a study for the first time, and when reloading the most recently used applicable workspace is loaded.
 [0293] Users can publish a study to a group or more users, which can also serve as a trigger for report generation and external system notifications.
 [0294] When opening a published workspace for the first time a private copy of the workspace is created with is also loaded on subsequent reloads. Published studies are immutable and can never be modified.
 [0295] Machine Learning with Medical Imaging
 [0296] With a cloud interface, it is now possible to aggregate statistics from multiple sources to come up with predictions using machine learning. These multiple sources can be the results generated by multiple people within an organization, or even multiple organizations scattered across the world. The statistics that can be aggregated can be medical imaging pixel data, medical imaging metadata (e.g. DICOM headers), and for example the electronic medical records of patients (EMRs). The learning can be applied at a user level, at an organization level, or even at a macro level (e.g. globally).
 [0297] In the case of trying to automatically quantify (e.g. annotate, measure, segment) medical images, there can be two different categories of deep learning, machine learning or artificial intelligence: For the medical imaging application, supervised learning is more appropriate because there is not sufficient data to learn from. In order to learn as effectively as possible, the cloud user interface has been tailored to allow users to add labels to the data in a structured fashion. For example, in the case of cardiovascular imaging, a user can make several measurements and label the measurements as they wish. Instead of allowing a completely user defined field, there is the option for a user to select a label from a predefined list that Arterys provides. By doing this, we can add labels to the data in a structured and automated fashion. Labeled data acts as the training data set

to feed into a machine learning algorithm (i.e. like a random forest or deep learning CNN or RNN) so that the algorithm can predict an outcome based on new unlabeled data. For example, one optional step in the user review process is for them to “publish” their workspace or state in a way that confirms that they are happy with the labels that they have added to the data. The “publish” mechanism can be an icon in the user interface that they click to “save”, or it can be the results that get sent to archive (for example to a hospital PACS server). There just needs to be a way to differentiate a user creating dummy measurements and annotations with true clinical measurements and annotations.

[0298] The benefit of a cloud interface is that every time a user makes any modification within the system interface to the suggestion provided, this modification then is saved and fed back into the machine learning labeled data. This creates a reinforcement learning loop that adds very valuable training data. The suggestions provided by the machine learning algorithm can be provided once when a user logs in or in real-time every time a user makes a modification during their session. For example, when a user identified a voxel in a medical image that is anatomy, all similar voxels can be identified in real-time in their session.

[0299] In the case of trying to predict the outcome of a particular treatment (and giving a resulting probability measure) or to predict which treatment choice is better suited for a particular patient, data from the EMR is critical. Having access to labeled medical device data (e.g. medical imaging, genomic data, wearables) is not sufficient in determining best treatment decisions. This data needs to be aggregated across all retrospective cases to offer a prediction to a new patient that has similar medical device data.

[0300] Machine learning can also be used for search in medical images. A user can type in a search field and find all images that for example has a particular type of disorder. A user can then verify that all the studies presented to them have this disorder and this data can then be fed back into the training dataset.

[0301] Picture and Video Service—

[0302] We want the user to be able to capture pictures and video of the current state of their workflow. These images and videos need to include both image data generated on our server and overlays rendered on the client browser. To accomplish this we have a node-webkit based video service that allows us to run both our client and server software stack within the same environment. We then restore the current state of the user’s workspace on the node-webkit environment and leverage the same compute nodes that were allocated for that user’s session. If a single picture is requested by the user the service simply takes a screenshot of the restored workspace and the resulting image file is returned. In the case of a video request the service takes a screenshot for each frame of the current workflow and compiles the screenshot images into a video file using a video encoder which is then returned. The returned images or video can be stored on the server or sent back to the client to be viewed.

[0303] Below is an example of a detailed software design for the picture and video service:

 Screenshots and Video Capture

 Requirements

- * screenshot should be a rendering of what the user currently sees in the viewport area
- * video can be an mpr cycled through time
- * video can be generated from a collection of key frames with in-betweens interpolating parameters that can be interpolated

-continued

- * video can be user interaction recording
- * output should contain everything that is on the viewport (image, webgl overlays, css overlays, ...)
- * screenshots and video frames should be full-quality

Design

Since we render everything on the client we need a client to generate the image.

Unfortunately, uploading a video would be prohibitive in the majority of network situations.

Hence, a screenshot/video service will run in the cluster that uses client rendering technology.

It will expose an interface over http to provide functionality defined in the requirements.

The service spins up node webkit processes on demand to render videos and screenshots as requests come in.

Upon receiving a request to render an image or collection of images, the service will launch a node webkit process and redirect it to signed URL for the user's worklist.

The node-webkit process will then load the study and inject the user's workspace

Next each frame will be rendered at full quality.

As a frame is rendered, node-webkit will perform an X11 screen capture and crop to the canvas viewport.

The image will be saved to disk.

Once all frames have been captured the service will return the screenshot, or in the case of a video,

a video will be encoded and returned.

Data Flow

- * user initiates request for screenshot or video.
- * webservice receives request
- * node-webkit process is started
- * node-webkit process opens a session, authenticated to load the required study
- * requested study is loaded
- * workspace in the request is injected into the study
- * upon completion of workspace load (including long running tasks like streamlines) it begins to render key frames
- * every frame is rendered full quality with no debounced image commands
- * as an image is rendered X11 screen grab (xwd) is executed on the window
- * the image is cropped to the viewport and saved to disk
- * if a video was requested encoding will run once images are generated
- * upon completion of all images an http response is sent back with the .png or .mp4
- * upon webservice receipt of the result, it will be saved in S3 and a reference saved to the database

Additional Tools and Optimizations

- * node-webkit requires webgl so the service will need to run on G2 instances
- * program 'xwd' in 'x11-apps' can capture window
- * ImageMagick 'convert' can convert xwd to png
- * ffmpeg can be used to generate .mp4 from a collection of .png

Details

Screenshots

+++++

Client message:

```
ws.emit('generate-screenshot', params);
```

params:

```
{
  workspace_id : 'workspace-id',
  id : 'app-id',
  study_id : 'study-id',
  upload_id : 'upload-id',
  window_width : 'browser_window_width',
  window_height : 'browser_window_height',
  hostname : window.location.hostname,
  port : window.location.port,
  pathname : '/app/xxx'
}
```

Video

+++++

-continued

```

Client message:
----
ws.emit('generate-screenshot', params);
----
params:
----
{
  // same as screenshot plus
  render_frames : [
    {
      orientation : [1,0,0,0,1,0,0,0,1],
      position : [0,0,0],
      time point : 1
      ...
    },
    {
      orientation : [0,1,0,1,0,0,0,0,1],
      position : [0,0,0],
      time point : 2
      ...
    }
  ],
  between : optional_number_of_frames_to_interpolate_between_frames
}
----
Webserver handler
++++

```

[0304] The message handler for 'generate-screenshot' attaches the current workspace to the args being sent to the webkit services

[0305] The webkit-client module is then used to send a request to one of the webkit services.

[0306] Once the response is received a record is inserted into the database and the image or video is stored.

Webkit-Client

[0307] +++++

[0308] The webkit-client module is responsible for routing a screenshot request to a node that can handle it.

[0309] The webkit-client subscribes to redis messages that are published by the currently running webkit nodes.

[0310] These messages include the existing instances of node-webkit that are running with the app-id they are running with.

[0311] When a request is received the webkit-client attempts to find a node that already has node-webkit running with the requested app-id.

[0312] Alternatively if a session has not been created yet it chooses the node with the least number of running sessions.

[0313] Once the node has been identified it sends the message over HTTPS to that host.

[0314] Arguments are sent in the body as JSON in a POST to the route '/webkit/execute'.

[0315] When the result returns a callback is invoked with the binary and a JSON blob containing the type (image/png or video/mp4), along with other useful information collected (e.g., timing information, size)

```

----
function execute(args, cb) {
  https.request('POST', '/webkit/execute', JSON.stringify(args),
    function(res) {

```

-continued

```

      cb(null, {
        binary : res,
        json : { type : headers['content-type'], info: {
          generation_start_time: <date>, generation_end_time: <date>}}
        });
      });
    }
  }
  ----

```

Webkit-Service

[0316] +++++

[0317] The webkit-service is a micro service that exposes an HTTPS interface to generate screenshots and videos.

[0318] The webkit-service listens only for POST requests at '/webkit/execute'.

[0319] Upon receiving a POST to '/webkit/execute' it creates a webkit-context and enqueues a request for a screenshot or video.

[0320] This module also takes care of authorizing the request that will be sent from node-webkit to the webserver by appending an auth_token associated with the special 'webkit-screenshot' user.

Webkit-Context

[0321] +++++

[0322] The webkit-context module is responsible for managing the node-webkit process that will run to generate a screenshot or video.

[0323] Upon creation a webkit-context creates a working directory to store intermediate results.

[0324] Next, it configures node-webkit by copying a simple 'index.html' and 'package.json' file into the working directory, and the 'args.json' containing the arguments passed into the context to render the screenshot/video.

[0325] Then node-webkit is started, and runs through the process of generating a screenshot.

[0326] When node-webkit exits, the webkit-context will look for the appropriate screenshot or video file to respond with.

[0327] Only one screenshot per app-id can run at a time.

[0328] A webkit-context registers itself in redis so that a webserver can route screenshot and video requests.

Node-Main

[0329] +++++++

[0330] The node-main module is the bridge module running in node-webkit.

[0331] When node-webkit starts it waits until the 'global.window' variable is defined, and then reads in the args.json file and starts executing the steps to generate a screenshot.

[0332] These arguments denote the width×height to make the window and where to redirect window.location.href to.

[0333] It assumes the redirect points to a website that will set global.window.io, which is an Arterys defined variable denoting the websocket connection.

[0334] Once the websocket connection has been made it invokes a 'load-study' command, and waits for 'load-workspace-complete'.

[0335] Once all commands that may have been invoked by restoring a workspace are finished node-main starts capturing images.

[0336] If 'args.json' contained the field 'render_frames' it iterates through each one generating an image.

[0337] The images are generated by invoking xwd to dump the Xwindow.

[0338] ImageMajick convert is then used to convert to a png and crop to the '.ar-content-body-canvas'.

[0339] If there was more than one image generated then ffmpeg is invoked to encode the collection of images into an h.264 encoded video.

[0340] When the screenshot or video has been created, node-webkit will exit cleanly.

[0341] Any error will cause node-webkit to exit with a non zero code, which will indicate to the webkit-context that the screenshot has failed.

[0342] The above described automated approaches remove the subjectivity in identifying anatomical structure and flow, which is endemic in conventional approaches, providing a high level of repeatability. This repeatability allows new uses for the MRI data. For example, MRI data for single patient may be reliably reviewed across different sessions for trends. Even more surprisingly, MRI data for a plurality of patients may be reliably reviewed for trends across a population or demographic.

[0343] The various embodiments described above can be combined to provide further embodiments. To the extent that they are not inconsistent with the specific teachings and definitions herein, all of the U.S. patents, U.S. patent application publications, U.S. patent applications, foreign patents, foreign patent applications and non-patent publications referred to in this specification and/or listed in the Application Data Sheet, including but not limited to U.S. Provisional Patent Application No. 61/571,908, filed Jul. 7, 2011; International Patent Application No. PCT/US2012/045575, filed Jul. 5, 2012; U.S. Provisional Patent Application No. 61/928,702, filed Jan. 17, 2014; International Patent Application No. PCT/US2015/011851, filed Jan. 16, 2015; U.S. patent application Ser. No. 15/112,130 filed Jul. 15, 2016;

U.S. Provisional Patent Application No. 62/260,565, filed Nov. 29, 2015; U.S. Provisional Patent Application No. 62/415,203 filed Oct. 31, 2016; and U.S. Provisional Patent Application No. 62/415,666 filed Nov. 1, 2016, are incorporated herein by reference, in their entirety. Aspects of the embodiments can be modified, if necessary, to employ systems, circuits and concepts of the various patents, applications and publications to provide yet further embodiments. [0344] These and other changes can be made to the embodiments in light of the above-detailed description. In general, in the following claims, the terms used should not be construed to limit the claims to the specific embodiments disclosed in the specification and the claims, but should be construed to include all possible embodiments along with the full scope of equivalents to which such claims are entitled. Accordingly, the claims are not limited by the disclosure.

1. A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically correct phase aliasing, the method comprising:

receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data comprising respective anatomical structure and velocity for each of a plurality of voxels;

for each of at least some of the plurality of voxels, identifying, by the at least one processor, sharp gradients in flow velocity near a velocity encoding parameter;

connecting, by the at least one processor, all of the voxels identified as having a sharp gradient to define an enclosed boundary;

determining, by the at least one processor, whether all of the voxels in the enclosed boundary are aliased; and responsive to determining that all of the voxels in the enclosed boundary are aliased:

adding, by the at least one processor, a multiple of the velocity encoding parameter to the velocity for each of the voxels in the enclosed boundary; or

subtracting, by the at least one processor, a multiple of the velocity encoding parameter to the velocity for each of the voxels in the enclosed boundary.

2. The method of claim 1, further comprising:

responsive to determining that not all of the voxels in the enclosed boundary are aliased, analyzing, by the at least one processor, the respective velocities of neighboring voxels in the enclosed boundary; and

modifying, by the at least one processor, the velocity of each of the voxels in the enclosed boundary based at least in part on the analyzed respective velocities of the neighboring voxels.

3. The method of claim 2 wherein modifying the velocity of each of the voxels in the enclosed boundary comprises modifying the velocity of each of the voxels by an amount determined to minimize the discontinuity across the neighboring voxels.

4. The method of claim 1 wherein determining whether all of the voxels in the enclosed boundary are aliased comprises determining whether there are any sharp gradients between neighboring voxels in the enclosed boundary.

5. A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system, the method comprising:

- receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data comprising respective anatomical structure and velocity information for each of a plurality of voxels;
- defining, by the at least one processor, a volume within the anatomical structure;
- minimizing, by the at least one processor, a divergence of a velocity field for at least some of the plurality of voxels that represent the volume within the anatomical structure; and
- correcting, by the at least one processor, velocity information for the at least some of the plurality of voxels which represent the volume within the anatomical structure based at least in part on the minimized divergence of the velocity field.
6. The method of claim 5 wherein correcting velocity information for the at least some of the plurality of voxels comprises correcting x velocity, y velocity and z velocity for the at least some of the plurality of voxels.
7. The method of claim 5 wherein minimizing a divergence of a velocity field comprises constructing a least squares divergence free approximation of the velocity field with a constraint of minimizing the divergence.
8. The method of claim 7 wherein minimizing a divergence of a velocity field comprises iteratively minimizing a residual divergence of the velocity field.
9. A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically correct phase aliasing, the method comprising:
- receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data comprising respective anatomical structure and velocity information for each of a plurality of voxels;
- analyzing, by the at least one processor, the set of MRI data to identify a time point which corresponds to a peak diastole of a cardiac cycle;
- analyzing, by the at least one processor, a variation of the velocity of at least some of the plurality of voxels over at least a portion of the cardiac cycle; and
- correcting, by the at least one processor, errors in the velocity of the at least some of the plurality of voxels based at least in part on the analysis of the variation of the velocity of at least some of the plurality of voxels over at least a portion of a cardiac cycle.
10. The method of claim 9 wherein analyzing a variation of the velocity of at least some of the plurality of voxels and correcting errors in the velocity of at least some of the plurality of voxels comprise:
- for each of the at least some of the plurality of voxels,
- tracking, by the at least one processor, a change in velocity between successive time points in the cardiac cycle;
- determining, by the at least one processor, that aliasing has occurred if the velocity varies by more than a magnitude of a velocity encoding parameter;
- responsive to determining that aliasing has occurred,
- incrementing, by the at least one processor, a wrap count if the velocity was reduced by more than the velocity encoding parameter; or
- decrementing, by the at least one processor, a wrap count if the velocity was increased by more than the velocity encoding parameter; and
- for each point in time, modifying, by the at least one processor, the velocity based at least in part on the accumulated wrap count for the voxel.
11. The method of claim 10 wherein modifying the velocity based at least in part on the accumulated wrap count for the voxel comprises increasing the velocity by an amount equal to product of the accumulated wrap count and two times the velocity encoding parameter.
12. The method of claim 10, further comprising:
- determining, by the at least one processor, whether the accumulated wrap count is zero over a cardiac cycle; and
- responsive to determining that the accumulated wrap count is not zero over a cardiac cycle, generating, by the at least one processor, an error signal.
13. The method of claim 12, further comprising:
- tracking, by the at least one processor, the number of voxels for which the accumulated wrap count is not zero over a cardiac cycle.
14. The method of claim 12, further comprising:
- correcting, by the at least one processor, errors in the velocity only for voxels for which the accumulated wrap count is not zero over a cardiac cycle.
15. The method of claim 9, further comprising:
- identifying, by the at least one processor, which of the plurality of voxels are likely to represent blood flow; and
- selecting, by the at least one processor, the at least some of the plurality of voxels for analysis based at least in part on the identification of which of the plurality of voxels are likely to represent blood flow.
16. A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically correct artifacts due to eddy currents, the method comprising:
- receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data comprising respective anatomical structure and velocity information for each of a plurality of voxels;
- receiving, by the at least one processor, an indication of which of the plurality of voxels represent static tissue;
- determining, by the at least one processor, at least one eddy current correction parameter based at least in part on the velocity information for the voxels which represent static tissue; and
- modifying, by the at least one processor, the velocity information for the plurality of voxels based at least in part the determined at least one eddy current correction parameter.
17. The method claim 16 wherein receiving an indication of which of the plurality of voxels represent static tissue comprises:
- filtering, by the at least one processor, the MRI data to mask regions of air;
- filtering, by the at least one processor, the MRI data to mask regions of blood flow; and
- filtering, by the at least one processor, the MRI data to mask regions of non-static tissue.

18. The method of claim **17** wherein filtering the MRI data to mask regions of air comprises masking regions with anatomy image values that are below a determined threshold.

19. The method of claim **18**, further comprising: analyzing, by the at least one processor, a histogram for anatomy image values to determine the determined threshold.

20. The method of claim **17**, further comprising: receiving, by the at least one processor, at least one filter parameter via a user interface communicatively coupled to the at least one processor.

21. A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically set an orientation of a multiplanar reconstruction, the method comprising:

receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data comprising respective anatomical structure and velocity information for each of a plurality of voxels;

presenting, by a display communicatively coupled to the at least one processor, the MRI data;

receiving, by the at least one processor, selection of a central region of blood flow;

determining, by the at least one processor, the direction of blood flow in the central region of blood flow; and

responsive to determining the direction of blood flow, adjusting, by the at least one processor, the orientation of the multiplanar reconstruction so that the multiplanar reconstruction is on a plane that is perpendicular to the determine direction of blood flow.

22. The method of claim **21** wherein determining the direction of blood flow in the central region of blood flow comprises:

determining a time point which corresponds to peak blood flow; and

determining the direction of blood flow at the time point which corresponds to peak blood flow.

23. A method of operation for use with a magnetic resonance imaging (MRI) based medical imaging system to automatically quantify blood flow in an anatomical volume, the method comprising:

receiving, by at least one processor, a set of MRI data representative of an anatomical structure, the set of MRI data comprising respective anatomical structure and velocity information for each of a plurality of voxels;

identifying, by the at least one processor, an anatomical volume in the set of MRI data which comprises a blood pool;

identifying, by the at least one processor, a plane on a landmark that is at least approximately perpendicular to the flow of blood in the anatomical volume;

generating, by the at least one processor, a contour defined by the intersection of the plane and the anatomical volume; and

determining, by the at least one processor, the total flow of blood through the voxels in the generated contour.

24. The method of claim **23** wherein determining the total flow of blood through the voxels in the generated contour comprises:

for each voxel in the generated contour, determining, by the at least one processor, a dot product of a normal vector of the plane and a velocity vector of the voxel; and

summing, by the at least one processor, the determined dot products of each of the voxels in the generated contour.

25. The method of claim **23** wherein the velocity information comprises a plurality of velocity components, and determining the total flow of blood through the voxels in the generated contour comprises:

for each velocity component, producing a velocity multiplanar reconstruction (MPR) in the plane of the generated contour, each of the velocity MPRs comprising a plurality of MPR pixels having pixel spacing at least similar to dimension of each of the voxels;

for each MPR pixel inside the generated contour, determining, by the at least one processor, a dot product of a normal vector of the plane and a velocity vector composed of the pixel values from the generated velocity MPRs; and

summing, by the at least one processor, the determined dot products of each of the MPR pixels in the generated contour.

26. A method of operation for use with a medical imaging system to automatically determine a volume of a particular region, the method comprising:

receiving, by at least one processor, a set of image data representative of an anatomical structure, the set of image data comprising respective anatomical structure information for each of a plurality of voxels;

receiving, by the at least one processor, an indication of a primary axis of a volume of interest;

generating, by the at least one processor, a plurality of spaced-apart slice planes along the primary axis of the volume of interest;

for each of the plurality of slice planes, generating, by the at least one processor, a closed contour which defines a boundary of the volume of interest at the slice plane;

generating, by the at least one processor, a three dimensional surface which connects the contours of all of the slice planes; and

determining, by the at least one processor, the volume of the three dimensional surface.

27. The method of claim **26** wherein the set of image data comprises respective anatomical structure information for each of a plurality of voxels at a plurality of time points, and the method comprises:

for each of the plurality of time points,

generating, by the at least one processor, a plurality of spaced-apart slice planes along the primary axis of the volume of interest;

for each of the plurality of slice planes, generating, by the at least one processor, a closed contour which defines a boundary of the volume of interest at the slice plane;

generating, by the at least one processor, a three dimensional surface which connects the contours of all of the slice planes; and

determining, by the at least one processor, the volume of the three dimensional surface.

28. The method of claim **26** wherein receiving an indication of a primary axis comprises receiving an indication of a primary axis that comprises one of a straight axis or a curved axis.

29. The method of claim **26** wherein the set of image data comprises respective anatomical structure information for each of a plurality of voxels at a plurality of time points, and the primary axis comprises a primary axis that moves over at least some of the plurality of time points.

30. The method of claim **26** wherein receiving an indication of a primary axis comprises receiving an indication of multiple axes.

31. A method of operation for use with a medical imaging system to automatically generate a connected mask, the method comprising:

receiving, by at least one processor, a set of image data representative of an anatomical structure, the set of image data comprising respective anatomical structure information for each of a plurality of locations in space; identifying, by the at least one processor, a seed location; identifying, by the at least one processor, an intensity value for the seed location; and

traversing, by the at least one processor, outward from the seed location to other locations to flood fill the connected mask according to flood fill criteria, the flood fill criteria including at least whether the intensity value of the each of the locations is within a specified threshold of the intensity value for the seed location.

32. The method of claim **31** wherein the received set of image data comprises three dimensional image data, the method further comprising:

generating, by the at least one processor, a two dimensional multiplanar reconstruction from the received image data.

33. The method of claim **31** wherein moving outward from the seed location to other locations to flood fill the mask comprises moving outward from the seed location to other locations to flood fill the mask according to flood fill criteria which includes connectivity criteria.

34. The method of claim **31** wherein moving outward from the seed location to other locations to flood fill the mask comprises moving outward from the seed location to other locations to flood fill the mask according to flood fill criteria, the flood fill criteria including at least one of a radius constraint or a number of flood fill steps constraint.

35. The method of claim **31**, further comprising:

generating, by the at least one processor, a contour defined by a border of the connected mask; and

generating, by the at least one processor, an approximation contour which is an approximation of the generated contour based at least in part on the generated contour.

36. The method of claim **35** wherein generating an approximation contour comprises generating an approximation contour using a spline which places control points at areas of high curvature.

37. The method of claim **35**, further comprising:

dilating, by the at least one processor, the approximation contour by a determined dilation factor to generate a dilated approximation contour.

38. A method of operation for use with a medical imaging system, the method comprising:

receiving, by at least one processor, a set of image data representative of an anatomical structure, the set of image data comprising respective anatomical structure information for each of a plurality of locations in three dimensional space;

tracking, by the at least one processor, the position and orientation of the anatomical structure at a plurality of time points;

at each of the plurality of time points, generating, by the at least one processor, a contour for the anatomical structure; and

determining, by the at least one processor, a measurement associated with the anatomical structure using the generated contours.

39. The method of claim **38** wherein determining a measurement associated with the anatomical structure using the generated contours comprises accounting for movement of the anatomical structure over at least some of the plurality of time points.

40. The method of claim **39** wherein accounting for movement of the anatomical structure comprises accounting for the linear and angular velocity of the anatomical structure over the at least some of the plurality of time points.

41. The method of claim **38** wherein determining a measurement associated with the anatomical structure using the generated contours comprises determining flow through the anatomical structure.

42. The method of claim **38**, further comprising:

presenting, by the at least one processor, a multiplanar reconstruction of the anatomical structure on a display at a plurality of time points, wherein the position and orientation of the multiplanar reconstruction tracks the generated contours over the plurality of time points.

43. The method of claim **38**, further comprising:

presenting, by the at least one processor, a multiplanar reconstruction of the anatomical structure on a display; and

presenting, by the at least one processor, the generated contours as semi-transparent if the multiplanar reconstruction is out of plane with the generated contours.

44. A processor-based device having at least one processor and at least one nontransitory processor-readable medium communicatively coupled to the at least one processor, and operable to perform any of the methods of claims **1** through **43**.

* * * * *

专利名称(译)	医学成像和医学成像信息的有效共享		
公开(公告)号	US20180256042A1	公开(公告)日	2018-09-13
申请号	US15/779447	申请日	2016-11-29
[标]申请(专利权)人(译)	阿特瑞斯公司		
申请(专利权)人(译)	ARTERYS INC.		
当前申请(专利权)人(译)	ARTERYS INC.		
[标]发明人	BECKERS FABIEN AXERIO CILIES JOHN TAERUM TORIN ARNI HSIAO ALBERT DE FRANCESCO GIOVANNI BIDULOCK DARRYL JUGDEV TRISTAN NEWTON ROBERT		
发明人	BECKERS, FABIEN AXERIO-CILIES, JOHN TAERUM, TORIN ARNI HSIAO, ALBERT DE FRANCESCO, GIOVANNI BIDULOCK, DARRYL JUGDEV, TRISTAN NEWTON, ROBERT		
IPC分类号	A61B5/026 A61B5/00 G16H30/40 G06T7/00 G06T7/269		
CPC分类号	A61B5/0263 A61B5/0013 A61B5/0044 A61B5/7203 G16H30/40 G06T7/0012 G06T7/269 A61B2576/023 G06T2207/10088 G06T2207/30104 A61B5/0022 A61B5/02 A61B5/02007 A61B5/02014 A61B5/021 A61B5/026 A61B5/055 A61B5/7207 A61B5/725 A61B5/7257 G01R33/5608 G01R33/56316 G06T2207/10076 G16H10/60 G16H30/20 H04L9/3213 H04L63/0209		
优先权	62/260565 2015-11-29 US 62/415203 2016-10-31 US		
外部链接	Espacenet USPTO		

摘要(译)

MRI图像处理和分析系统可以基于所识别的结构识别MRI流数据中的结构的实例，例如，一致性，导出轮廓和/或临床标记。该系统可以远离一个或多个MRI采集系统，并执行：对MRI数据集执行错误检测和/或校正（例如，相位误差校正，相位混叠，信号展开和/或其他伪像）；分割；流动的可视化（例如，速度，动脉与静脉血流，分流）叠加在解剖结构上，量化；验证；和/或产生患者特异性4-D流程协议。即使对于复杂或大型医学成像数据集和元数据，异步命令和成像管道也允许以及时和安全的方式进行远程图像处理和分析。

