

MODEL PARAMETER DETERMINATION USING A PREDICTIVE MODEL

CROSS-REFERENCE TO RELATED APPLICATION

This application claims priority under 35 U.S.C. § 120 to: U.S. Non-Provisional Application Ser. No. 16/277,939 filed on February 15, 2019, the contents of which are hereby incorporated by reference in their entirety.

BACKGROUND

Field

[001] The described embodiments relate generally to determining model parameters associated with a sample for use in a forward model that simulates response physics of the sample to an excitation.

Related Art

[002] Many non-invasive characterization techniques are available for determining one or more physical parameters of a sample. For example, magnetic properties can be studied using magnetic resonance or MR (which is often referred to as ‘nuclear magnetic resonance’ or NMR), a physical phenomenon in which nuclei in a magnetic field absorb and re-emit electromagnetic radiation. Moreover, density variations and short or long-range periodic structures in solid or rigid materials can be studied using characterization techniques such as x-ray imaging, x-ray diffraction, computed tomography, neutron diffraction or electron microscopy, in which electromagnetic waves or energetic particles having small de Broglie wavelengths are absorbed or scattered by the sample. Furthermore, density variations and motion in soft materials or fluids can be studied using ultrasound imaging, in which ultrasonic waves are transmitted and reflected in the sample.

[003] In each of these and other non-invasive characterization techniques, one or more external excitation (such as a flux of particles or incident radiation, static or time-varying scalar fields, and/or static or time-varying vector fields) are applied to the sample, and a resulting response of the sample, in the form a physical phenomenon, is measured to, directly or indirectly, determine the one or more physical parameters. As an example, in MR magnetic nuclear spins may be partially aligned (or polarized) in an applied external DC magnetic field. These nuclear spins may precess or rotate around the direction of the external

magnetic field at an angular frequency (which is sometimes referred to as the ‘Larmor frequency’) given by the product of a gyromagnetic ratio of a type of nuclei and the magnitude or strength of the external magnetic field. By applying a perturbation to the polarized nuclear spins, such as one or more radio-frequency (RF) pulses (and, more generally, electro-magnetic pulses) having pulse widths corresponding to the angular frequency and at a right-angle or perpendicular to the direction of the external magnetic field, the polarization of the nuclear spins can be transiently changed. The resulting dynamic response of the nuclear spins (such as the time-varying total magnetization) can provide information about the physical and material properties of a sample, such as one or more physical parameters associated with the sample.

[004] Moreover, in general each of the characterization techniques may allow one or more physical parameters to be determined in small volumes or voxels in a sample, which can be represented using a tensor. Using magnetic resonance imaging (MRI) as an example, the dependence of the angular frequency of precession of nuclear spins (such as protons or the isotope ^1H) on the magnitude of the external magnetic field can be used to determine images of three-dimensional (3D) or anatomical structure and/or the chemical composition of different materials or types of tissue. In particular, by applying a non-uniform or spatially varying magnetic field to a sample, the resulting variation in the angular frequency of precession of ^1H spins is typically used to spatially localize the measured dynamic response of the ^1H spins to voxels, which can be used to generate images, such as of the internal anatomy of a patient.

[005] However, the characterization of the physical properties of a sample is often time-consuming, complicated and expensive. For example, acquiring MR images in MRI with high-spatial resolution (i.e., small voxels sizes) often involves a large number of measurements (which are sometimes referred to as ‘scans’) to be performed for time durations that are longer than the relaxation times of the ^1H spins in different types of tissue in a patient. Moreover, in order to achieve high-spatial resolution, a large homogenous external magnetic field is usually used during MRI. The external magnetic field is typically generated using a superconducting magnetic having a toroidal shape with a narrow bore, which can feel confining to many patients. Furthermore, Fourier transform techniques may

be used to facilitate image reconstruction, at the cost of constraints on the RF pulse sequences and, thus, the MR scan time.

[006] The combination of long MR scan times and, in the case of MRI, the confining environment of the magnet bore can degrade the user experience. In addition, long MR scan times reduce throughput, thereby increasing the cost of performing the characterization. These types of problems can constrain or limit the use of many characterization techniques.

[007] One approach for addressing these problems is to use simulations of the response physics of a sample to one or more excitations to determine information, such as the one or more physical parameters. For example, using model parameters at the voxel level and a forward model based on one or more differential equations that describe a physical phenomenon, a computer can simulate the response physics of the sample as an output of the forward model using information specifying the one or more excitations as an input to the forward model.

[008] However, this approach often replaces the problems of having a large number of MR scans and long MR scan times, with the problems associated with accurately determining the model parameters at the voxel level. For example, the model parameters are typically determined by iteratively applying one or more excitations, performing measurements and then solving an inverse problem of using the measurements to compute the corresponding model parameters until a desired accuracy of the simulated response physics is achieved (which is sometimes referred to as an ‘iterative approach’). In general, it can be difficult, time-consuming and expensive to determine the model parameters using these existing techniques, which may constrain or limit the use of simulations of the response physics to characterize a sample.

SUMMARY

[009] A system that determines model parameters associated with a sample is described. This system includes: a source configured to provide an excitation; a measurement device that performs measurements; a processor that executes program instructions; and memory that stores the program instructions. During operation, the system may apply, to the sample, the excitation using the source, where the excitation has at least a wavelength and an intensity or a flux, and where the excitation includes one of: electromagnetic radiation, a radio-frequency wave, a particle beam, a sound wave, a magnetic field, or an electric field. Then,

the system may measure, using a measurement device, a response associated with the sample to the excitation. Moreover, the system may compute, using the measured response and information specifying the excitation as inputs to a predetermined predictive model, model parameters on a voxel-by-voxel basis in a forward model with multiple voxels that represent the sample. The forward model may simulate response physics occurring within the sample to a given excitation with a given wavelength and a given intensity or a given flux, that are selected from a range of measurement conditions that includes the excitation, the wavelength and the intensity or the flux, and at least a different wavelength and a at least a different intensity or a different flux. Furthermore, the forward model may be a function of the excitation, the model parameters of the multiple voxels, and differential or phenomenological equations that approximates the response physics. Next, the system may determine, using the processor, an accuracy of the model parameters by comparing at least the measured response and a calculated predicted value of the response using the forward model, the model parameters and the excitation. Additionally, when the accuracy exceeds a predefined value, the system may provide the model parameters as an output to a user, to another electronic device, to a display and/or to the memory.

[010] For example, the excitation may include at least one of: an electromagnetic beam in an x-ray band of wavelengths, a neutron beam, an electron beam, an electromagnetic beam in an optical band of wavelengths, an electromagnetic beam in an infrared band of wavelengths, a sound wave in an ultrasound band of wavelengths, a proton beam, an electric field associated with an impedance measurement device, a radio-frequency wave associated with a magnetic-resonance apparatus, and/or a magnetic field associated with a susceptibility measurement device.

[011] Moreover, the measurement device may include at least one of: an x-ray detector, a neutron detector, an electron detector, an optical detector, an infrared detector, an ultrasound detector, a proton detector, the magnetic-resonance apparatus, the impedance measurement device and/or the susceptibility measurement device.

[012] Note that the predetermined predictive model may include a machine-learning model and/or a neural network. In some embodiments, the predetermined predictive model includes a personalized predictive model that corresponds to an individual.

[013] Furthermore, when the accuracy exceeds the predefined value, the model parameters may be computed in a single pass without further iteration. Consequently, the model parameters having an accuracy exceeding the predefined value may be computed using fewer iterations with the predetermined predictive model than in the iterative approach without the predetermined predictive model.

[014] Alternatively, when the accuracy is less than the predefined value, the system may: calculate, using information specifying the excitation and the accuracy as inputs to a second predetermined predictive model, a revised excitation that has at least a revised wavelength, a revised intensity or a revised flux; and repeat the applying, measuring, computing and determining operations with the revised excitation instead of the excitation. Note that the second predetermined predictive model may include a machine-learning model and/or a neural network.

[015] In some embodiments, when the accuracy is less than the predefined value, the system may: add the excitation and the measured response to a training dataset; and determine, using the training dataset, a revised instance of the predictive model.

[016] Additionally, the system may classify or segment one or more anatomical structures in the sample using the model parameters and a third predictive model. For example, the third predetermined predictive model may include a machine-learning model and/or a neural network.

[017] Moreover, the system may train the predictive model using a simulated dataset computed using the forward model, a range of model parameters and a range of excitations.

[018] Furthermore, the measured response may include a time-domain response of the sample and may be other than an image.

[019] Another embodiment provides a computer-readable storage medium for use with the system. This computer-readable storage medium includes a program module that, when executed by the system, causes the system to perform at least some of the aforementioned operations.

[020] Another embodiment provides a method for determining the model parameters associated with the sample. This method includes at least some of the aforementioned operations performed by the system.

[021] This Summary is provided for purposes of illustrating some exemplary embodiments, so as to provide a basic understanding of some aspects of the subject matter described herein. Accordingly, it will be appreciated that the above-described features are simply examples and should not be construed to narrow the scope or spirit of the subject matter described herein in any way. Other features, aspects, and advantages of the subject matter described herein will become apparent from the following Detailed Description, Figures, and Claims.

BRIEF DESCRIPTION OF THE FIGURES

[022] FIG. 1 is a block diagram illustrating an example of a system in accordance with an embodiment of the present disclosure.

[023] FIG. 2 is a flow diagram illustrating an example of a method for determining model parameters associated with a sample in accordance with an embodiment of the present disclosure.

[024] FIG. 3 is a drawing illustrating an example of communication among components in the system in FIG. 1 in accordance with an embodiment of the present disclosure.

[025] FIG. 4 is a drawing illustrating an example of a machine-learning model in accordance with an embodiment of the present disclosure.

[026] FIG. 5 is a drawing illustrating an example of a neural model in accordance with an embodiment of the present disclosure.

[027] FIG. 6 is a drawing illustrating an example of classification or segmentation of one or more anatomical structures in the sample in accordance with an embodiment of the present disclosure.

[028] FIG. 7 is a block diagram illustrating an example of an electronic device in accordance with an embodiment of the present disclosure.

[029] FIG. 8 is a drawing illustrating an example of a data structure that is used by the electronic device of FIG. 7 in accordance with an embodiment of the present disclosure.

[030] Note that like reference numerals refer to corresponding parts throughout the drawings. Moreover, multiple instances of the same part are designated by a common prefix separated from an instance number by a dash.

DETAILED DESCRIPTION

[031] A system that determines model parameters associated with a sample is described. During operation, the system may apply, to the sample, the excitation using the source.

Then, the system may measure, using a measurement device, a response associated with the sample to the excitation. Moreover, the system may compute, using the measured response and information specifying the excitation as inputs to a predetermined predictive model, model parameters on a voxel-by-voxel basis in a forward model with multiple voxels that represent the sample. The forward model may simulate response physics occurring within the sample to a given excitation. Furthermore, the forward model may be a function of the excitation, the model parameters of the multiple voxels, and differential or phenomenological equations that approximates the response physics. Next, the system may determine, using the processor, an accuracy of the model parameters by comparing at least the measured response and a calculated predicted value of the response using the forward model, the model parameters and the excitation. Additionally, when the accuracy exceeds a predefined value, the system may provide the model parameters as an output to a user, to another electronic device, to a display and/or to a memory.

[032] By determining the model parameters for voxels in the sample (which is sometimes referred to as ‘tensor field mapping’ or TFM, because the parameters in the voxels can be represented by a hybrid tensor as opposed to a true tensor for a vector field), this computation technique may reduce or eliminate the need for iterative measurements and adaptation when determining the model parameters. Consequently, the computation technique may significantly reduce the use of system resources (such as processor time, memory, etc.) when determining the model parameters. Moreover, if the accuracy is insufficient (such as when the accuracy is less than the predefined value), the computation technique may be used to guide a modification to the excitation to facilitate rapid convergence on the model parameters with the desired accuracy. Furthermore, by providing a forward model that predicts a physical phenomenon based on the determined model parameters for a range of excitation values or intensities, the computation technique may facilitate rapid and accurate characterization of a sample (such as the determination of one or more physical parameters of a sample). Therefore, the computation technique can be used to dynamically adapt or modify the excitation used in the measurements and/or may facilitate improved sample characterization.

[033] These capabilities may result in shorter MR scan or measurement times, increased throughput and, thus, reduced measurement cost, an improved user experience (such as by

reducing the amount of time people spend in the confining environment of a magnet bore in an MR scanner), and increased use of characterization techniques. In addition, the computation technique may facilitate quantitative analysis of measurements, which may improve the accuracy, may reduce errors and, thus, may improve the health and well-being of people.

[034] In general, the computation technique may be used in conjunction with a variety of characterization techniques and forward models that quantitatively simulate the response physics occurring within the sample to a given excitation. For example, the characterization technique may involve: x-ray measurements (such as x-ray imaging, x-ray diffraction or computed tomography), neutron measurements (neutron diffraction), electron measurements (such as electron microscopy or electron spin resonance), optical measurements (such as optical imaging or optical spectroscopy that determines a complex index of refraction at one or more wavelengths), infrared measurements (such as infrared imaging or infrared spectroscopy that determines a complex index of refraction at one or more wavelengths), ultrasound measurements (such as ultrasound imaging), proton measurements (such as proton scattering), MR measurements or an MR technique (such as MRI, MR spectroscopy or MRS with one or more types of nuclei, magnetic resonance spectral imaging or MRSI, MR elastography or MRE, MR thermometry or MRT, magnetic-field relaxometry, diffusion-tensor imaging and/or another MR technique, e.g., functional MRI, metabolic imaging, molecular imaging, blood-flow imaging, etc.), impedance measurements (such as electrical impedance at DC and/or an AC frequency) and/or susceptibility measurements (such as magnetic susceptibility at DC and/or an AC frequency). Therefore, the excitation may include at least one of: an electromagnetic beam in an x-ray band of wavelengths (such as between 0.01 and 10 nm), a neutron beam, an electron beam, an electromagnetic beam in an optical band of wavelengths (such as between 300 and 800 nm), an electromagnetic beam in an infrared band of wavelengths (such as between 700 nm and 1 mm), a sound wave in an ultrasound band of wavelengths (such as between 0.2 and 1.9 mm), a proton beam, an electric field associated with an impedance measurement device, a radio-frequency wave associated with an MR apparatus or scanner, and/or a magnetic field associated with a susceptibility measurement device. However, another non-invasive characterization technique (such as positron emission spectroscopy), an integrated therapy (such as proton beam therapy or

proton implantation, radiation therapy, magnetically guided nano particles, etc.) and/or a different range of wavelengths (such as ultraviolet wavelengths between 10 and 400 nm) may be used. In general, the computation technique may be used with a wide variety of excitation may be used to ‘excite’ a region of space as long as there is a forward model that describes the response physics for these excitations. In the discussion that follows, an MR technique is used as an illustrative example of a characterization technique.

[035] Note that the sample may include an organic material or an inorganic material. For example, the sample may include: an inanimate (i.e., non-biological) sample, a biological lifeform (such as a person or an animal, i.e., an in-vivo sample), or a tissue sample from an animal or a person (i.e., a portion of the animal or the person). In some embodiments, the tissue sample was previously removed from the animal or the person. Therefore, the tissue sample may be a pathology sample (such as a biopsy sample), which may be formalin fixed-paraffin embedded. In the discussion that follows, the sample is a person or an individual, which is used as an illustrative example.

[036] We now describe embodiments of a system. FIG. 1 presents a block diagram illustrating an example of a system 100. In system 100, a source 110 selectively provides an excitation to a sample 112, and a measurement device 114 selectively performs measurements on sample 112 to measure a response of sample 112 to the excitation. Moreover, system 100 includes a computer 116. As described further below with reference to FIG. 7, computer 116 may include subsystems, such as a processing subsystem, a memory subsystem and a networking subsystem. For example, the processing subsystem may include a processor that executes program instructions, the memory subsystem may include a memory that stores the program instructions, and networking subsystem may include an interface that communicates instructions or commands to source 110 and measurement device 114 (such as one or more sensors), that receives measurements from measurement device 114, and that selectively provides determined model parameters.

[037] During operation, a communication engine (or module) 120 in computer 116 may provide, via a network 118 (such as one or more wired and/or wireless links or interconnects), an instruction or a command to source 110, which may cause source 110 to apply, to sample 112, the excitation. This excitation may have at least a wavelength and an intensity or a flux. For example, the excitation may include: electromagnetic radiation, a

radio-frequency wave, a particle beam, a sound wave, a magnetic field, and/or an electric field.

[038] In some embodiments, the excitation may include an external magnetic field that polarizes one or more types of nuclei in sample 112, an optional gradient in the magnetic field, and/or a radio-frequency (RF) pulse sequence (which are sometimes referred to as ‘measurement conditions’ or ‘scanning instructions’). Thus, source 110 may include a magnet that applies the external magnetic field, an optional gradient coil that applies the optional gradient, and/or an RF coil that applies the RF pulse sequence.

[039] Then, communication engine 120 may provide, via network 118, an instruction or a command to measurement device 114, which may cause measurement device 114 to perform measurements of the response of at least a portion of sample 112 to the excitation. Moreover, measurement device 114 may provide, via network 118, the measurement results to communication engine 120. Note that measurement device 114 may include: an x-ray detector, a neutron detector, an electron detector, an optical detector, an infrared detector, an ultrasound detector, a proton detector, an MR apparatus or scanner, the impedance measurement device (such as a gel-covered table in an MR apparatus or scanner) and/or the susceptibility measurement device.

[040] In some embodiments, measurement device 114 may include one or more RF pickup coils or another magnetic sensor (such as a magnetometer, a superconducting quantum interference device, opto-electronics, etc.) that measure time-varying or time-domain electrical signals corresponding to the dynamic behavior of nuclear spins in the one or more types of nuclei or at least an average component of the magnetization corresponding to the aggregate dynamic behavior of the nuclear spins (which is sometimes referred to as a ‘magnetic response’) of at least the portion of sample 112. For example, measurement device 114 may measure the transverse magnetization of at least a portion of sample 112 as it precesses in the xy plane.

[041] Note that the measurements provided by measurement device 114 may be other than or different from an image. For example, the measurements may be other than MRI results. For example, the measurements may include or may correspond to (such as one or more components of) a free-induction-decay of the nuclear spins in sample 112. Consequently, in some embodiments the measurements may not involve performing a Fourier transform on the

measured electrical signals (and, thus, may not be performed in k-space and may not involve pattern matching in k-space, such as MR fingerprinting). However, in general, the measurements may be specified in the time domain and/or the frequency domain. Therefore, in some embodiments, a variety of signal processing (such as filtering, image processing, etc.), noise cancellation and transformation techniques (such as a discrete Fourier transform, a Z transform, a discrete cosine transform, data compression, etc.) may be performed on the measurements.

[042] After receiving the measurements, analysis engine (or module) 122 in computer 116 may analyze the measurements. This analysis may involve determining a (possibly time-varying) 3D position of sample 112 relative to measurement device 114 (which is sometimes referred to as ‘3D registration information’). For example, the aligning may involve performing point-set registration, such as with reference markers at known spatial locations. The registration may use a global or a local positioning system to determine changes in the position of sample 112 relative to measurement device 114. Alternatively or additionally, the registration may be based at least in part on variation in the Larmor frequency and the predetermined spatial magnetic-field inhomogeneity or variation in the magnetic field of source 110 and/or measurement device 112 (such as an MR apparatus or scanner). In some embodiments, the analysis involves aligning the voxels based at least in part on the registration information with desired voxel locations, and/or resampling and/or interpolating the measured signals to different voxel locations, which may facilitate subsequent comparisons with previous measurements or results.

[043] Moreover, analysis engine 122 may use the measurements to determine model parameters for a forward model with multiple voxels that represent sample 112, and that simulates the response physics occurring in sample 112 to a given excitation in a range of possible excitations (i.e., the forward model may be more general than one that determines the predicted response to a particular or a specific excitation). Notably, with the appropriate model parameters for the voxels in sample 112, analysis engine 122 may use the forward model to accurately and quantitatively simulate or calculate a predicted response of sample 112 to the excitation (such as a predicted component of the magnetization). Note that the forward model may be based at least in part on or may use one or more differential equations or one or more phenomenological equations that approximates the response physics of

sample 112 on a voxel-by-voxel basis. For example, the forward model may be based at least in part on or may use the Bloch equations, the Bloch-Torrey equations (thus, the forward model may include simulations of dynamics, such as motion associated with: respiration, a heartbeat, blood flow, mechanical motion, etc.), full Liouvillian computations (such as a Liouville supermatrix of interactions between two or more elements), a full Hamiltonian, Maxwell's equations (e.g., the forward model may calculate magnetic and electrical properties of sample 112), thermal diffusion equations, the Pennes equations, and/or another simulation technique that represents the physics of a response of sample 112 to a type of excitation. Because in some embodiments the assumptions underlying the Bloch equations are invalid (such as the parallel and antiparallel components of the magnetization are coupled, e.g., when the state of the magnetization is not reset prior to an RF pulse sequence), additional error terms may be added to the Bloch equations. Therefore, the forward model may be able to compute a dynamic (e.g., time-varying) state of sample 112 in response to an arbitrary excitation in a range of possible excitations or excitation values.

[044] In some analysis approaches, computer 116 may determine the model parameters by solving an inverse problem by iteratively modifying the model parameters associated with the voxels in the forward model until a difference between the predicted response and the measured dynamic magnetic response is less than a predefined value (such as 0.1, 1, 5 or 10%). (Note that 'an inverse problem' starts with one or more result(s) or output(s) and then calculates the inputs or causes. This is the inverse of a 'forward problem,' which starts with the inputs and then calculates the one or more results or the outputs.) However, in this 'iterative approach,' source 110 may repeatedly apply different excitations, and measurement device 114 may repeatedly perform corresponding measurements. Consequently, the iterative approach may be time-consuming, expensive and complicated. Thus, the iterative approach may consume significant resources in system 100 until the appropriate model parameters are determined.

[045] As described further below with reference to FIGs. 2-5, in order to address these problems, in the computation technique analysis engine 122 may use one or more predetermined or pretrained predictive models (such as a machine-learning model or a neural network, which may be specific to a particular sample or an individual, e.g., the predictive model may be a personalized predictive model) to, at least in part, compute the model

parameters on a voxel-by-voxel basis. For example, analysis engine 122 may use the measurements and information specifying the excitation as inputs to a predictive model, which provides, as an output, the model parameters associated with the voxels. Therefore, the predictive model may be trained on or may incorporate model-parameter information based at least in part on measurements or measurement results. In some embodiments, the predictive model may correct the measurements for extrinsic characteristics or a signature of a specific source 110 and/or measurement device 114 (such as RF noise or spatial magnetic-field inhomogeneity) and/or a particular excitation or measurement condition, so that the determined model parameters are intrinsic to sample 112 at a particular time when the measurements were performed.

[046] Note that the model parameters may include: a spin-lattice relaxation time T_1 (which is the time constant associated with the loss of signal intensity as components of the nuclear-spin magnetization vector of a type of nuclei relax to be parallel with the direction of an external magnetic field), a spin-spin relaxation time T_2 (which is the time constant associated with broadening of the signal during relaxation of components of the nuclear-spin magnetization vector of a type of nuclei perpendicular to the direction of the external magnetic field), an adjusted spin-spin relaxation time T_2^* , proton or nuclei density (and, more generally, the densities of one or more type of nuclei), diffusion (such as components in a diffusion tensor), velocity/flow, temperature, off-resonance frequency, electrical conductivity or a dielectric constant, and/or a magnetic susceptibility or permittivity.

[047] If a subsequent simulation using these model parameters provided by the predictive model, the forward model and one or more excitations of one or more predicted responses of sample 112 (such as a simulated or predicted MR signal) agrees with the corresponding measurements (such as a difference between a predicted response and a measurement is less than a predefined value, e.g., 0.1, 1, 5 or 10%, or alternatively when an accuracy exceeds a predefined value), results engine (or module) 124 in computer 116 may provide the determined model parameters, such as by providing an output to a user, to another electronic device, to a display and/or to the memory. In some embodiments, results engine 124 may output a tensor field map for sample 112 with model parameters for 3 spatial x one temporal x up to N measurement dimensions, where each measurement may be a vector or scalar quantity.

[048] Thus, when the accuracy exceeds the predefined value (such as 90, 95, 99 or 99.9%), the model parameters may be computed in a single pass without further iteration. Consequently, the model parameters having an accuracy exceeding the predefined value may be computed using fewer (or no) iterations with the predetermined predictive model (and, thus, more rapidly) than in the iterative approach without the predetermined predictive model.

[049] Alternatively, when the accuracy is less than the predefined value, computer 116 may perform one or more iterations in which one or more different, modified or revised excitations (such as a different RF pulse sequence) are applied to sample 112 by source 114, and one or more corresponding additional measurements are performed by measurement device 114. These one or more additional measurements may be used by computer 116 to determine the model parameters with an accuracy less than the predefined value.

[050] For example, analysis engine 122 may use a second predetermined predictive model (such as a second machine-learning model or a second neural network) to determine a revised excitation. Notably, using information specifying the excitation and the accuracy as inputs, the second predictive model may output the revised excitation. Then, system 100 may repeat the applying, measuring, computing and determining operations with the revised excitation instead of the excitation. Therefore, the second predictive model may be trained on or may incorporate excitation information based at least in part on remaining differences between the predicted response and the measurement in order to reduce or eliminate the remaining differences in one or more subsequent iterations of the operations performed by system 100. In some embodiments, the second predictive model may revise a sampling frequency, a characterization technique, etc. to determine additional information that allows the determination of the model parameters using the first predictive model to converge (i.e., to have an accuracy less than the predefined value). Stated differently, the next perturbation or disturbance may be chosen to minimize the error or the difference across the hyper-dimensional space.

[051] In some embodiments, when the accuracy is less than the predefined value, training engine (or module) 126 in computer 116 may: add the excitation and the measured response to a training dataset; and determine, using the training dataset, a revised instance of the predictive model for subsequent use in determining the model parameters. Thus, the

measurements performed by system 100 may be selectively used in an adaptive learning technique to improve the predictive model and, therefore, the determined model parameters for a range of excitations (such as different values of the wavelength and the intensity or the flux).

[052] Using the model parameters and the forward model, analysis engine 122 may simulate or predict a response of sample 112 to an arbitrary excitation, such as an arbitrary external magnetic field strength or direction (such as 0 T, 6.5 mT, 1.5 T, 3 T, 4.7 T, 9.4 T, and/or 15 T, or a time-varying direction, e.g., a slowly rotating external magnetic field), an arbitrary optional gradient, an arbitrary pulse sequence, an arbitrary magnetic state or condition (e.g., in which the magnetization or polarization of sample 112 is not returned to, been reset to or re-magnetized to an initial state prior to a measurement), etc. Therefore, the model parameters and the forward model may be used to facilitate fast and more accurate measurements, such as: soft-tissue measurements, morphological studies, chemical-shift measurements, magnetization-transfer measurements, MRS, measurements of one or more types of nuclei, Overhauser measurements, and/or functional imaging. For example, in embodiments where computer 116 determines the model parameters concurrently with measurements performed on sample 112 by source 110 and measurement device 114 (i.e., in real time), system 100 may rapidly characterize one or more physical parameters of sample 112 (at the voxel level or on average) on time scales smaller than T_1 or T_2 in an arbitrary type of tissue. This capability may allow system 100 to perform initial measurements to determine the model parameters, and then to use the determined model parameters to simulate or predict MR signals to complete or fill in ongoing measurements being performed by system 100, so that the results can be obtained more rapidly (and, thus, with a shorter MR scan time). Note that, in some embodiments, system 100 may determine the results (such as detecting an anomaly or a change in sample 112) based at least in part on quantitative comparisons of previous results obtained on sample 112, such as stored model parameters for the voxels in sample 112 that were determined during a previous MR scan(s) of sample 112. Such comparisons may be facilitated by 3D registration information that allows the voxel positions in sample 112 at different times to be aligned. In some embodiments, the results are based at least in part on a physician's instructions, medical lab test results (e.g., a blood test, urine-sample testing, biopsies, genetic or genomic testing, etc.), an individual's medical

history, the individual's family history, quantitative tensor field maps with voxel-dependent multi-dimensional data for sample 112 or other samples, impedance of sample 112, a hydration level of sample 112 and/or other inputs.

[053] Furthermore, as described further below with reference to FIG. 6, in some embodiments analysis engine 122 may classify or segment one or more anatomical structures in sample 112 using the determined model parameters and a third predetermined predictive model (such as a third machine-learning model and/or a third neural network). For example, using the simulated or predicted response of sample 112 at the voxel level or the determined model parameters at the voxel level, the third predictive model may output the locations of different anatomical structures and/or may output classifications of different voxels (such as a type of organ, whether they are associated with a particular disease state, e.g., a type of cancer, a stage of cancer, etc.). Therefore, in some embodiments, the third predictive model may be trained on or may incorporate classification or segmentation information based at least in part on variation in the model parameters across boundaries between different voxels (such as discontinuous changes). This capability may allow analysis engine 122 to identify different anatomical structures (which may assist in the determination of the model parameters) and/or to diagnose or to make a diagnosis recommendation about a medical condition or a disease state. In some embodiments, the classification or segmentation is performed prior to, concurrently or following the determination of the model parameters.

[054] In some embodiments, training engine 126 may have, at least in part, trained the predictive model, the second predictive model and/or the third predictive model using a simulated dataset. For example, training engine 126 may have generated the simulated dataset using the forward model, a range of model parameters and a range of excitations. In this way, simulated data may be used to accelerate training of one or more predictive models.

[055] Notably, because the computation technique may capture all relevant information during the measurements (such as an MR scan), the forward model can be used in an off-line mode to curate an extensive, labeled dataset that includes a large number of possible scenarios (such as different measurement conditions). This database can then be used to train predictive models. This capability may address the difficulty in obtaining MR data that is accurately labeled, reproducible, and artifact-free.

[056] In conjunction with the generated dataset, one or more predictive models can be used to select regularization that accelerates the initial data acquisition and/or denoising.

Moreover, the one or more predictive models can also be used to accelerate simulations or reconstruction using the forward model. For example, a predictive model can provide initial model parameters for use in the forward model, which may reduce the number of iterations required for the measurements and the simulations to converge on a solution that has an accuracy exceeding the predefined value. Thus, if the initial model parameters result in predicted response that are very different from the measurements, this can be feedback into the subsequent measurements and simulations to improve the model parameters and, thus, the predicted response.

[057] Furthermore, if there is a portion of the model-parameter space that is not covered by the predictive model(s), new data points can be accurately generated and labeled to train the predictive model(s). Additionally, the predictive model(s) may be trained based on different metrics corresponding to different applications. For example, the predictive model(s) may be training to optimize the excitations used in difference scenarios (such as fast scanning for asymptomatic population, high accuracy for specific tissue properties, robustness to variations in the signal-to-noise ratio, different hardware imperfections, etc.).

[058] In some embodiments, analysis engine 122 may run a neural network that determines first model parameters based at least in part on measured or simulated data and may performs brute-force nonlinear numerical calculations to solve an inverse problem using the measured or the simulated data to determine second model parameters. The difference between the first and the second model parameters from these two ‘inverse solvers’ may be used as the error in the neural-network-based approach. This approach may allow the neural network to learn because the numerical approach may be able to give real-time feedback to the neural network and to back propagate/update the weights in the neural network. This hybrid approach would still not require or need a priori training, but would be able to leverage the pattern-matching benefits of large neural networks with the determinism and accuracy of simulation/numerical techniques to solve the inverse problem. The hybrid approach may assist the neural network when it has an input unlike any of the examples used to train it. Similarly, the hybrid approach may be used to go directly from time-domain measurement to the model-

parameterized output (i.e. the inverse problem outputs). In some embodiments, the hybrid approach is implemented using a generative adversarial network (GAN).

[059] Note that, in some embodiments, the forward model may be independent of a particular MR apparatus or scanner. Instead, the forward model may be, e.g., specific to an individual. The predicted response computed using the forward model may be adjusted to include characteristics or a signature of a particular MR apparatus or scanner, such as magnetic-field inhomogeneity or spatial variation in the magnetic field, RF noise, a particular RF pickup coil or another magnetic sensor, variation in the characteristics or the signature with the external magnetic-field strength or the measurement conditions (such as the voxel size), geographic location, time (due to, e.g., magnetic storms), etc. Thus, the predicted response may be machine-specific.

[060] While the preceding discussion illustrated the computation technique using a single predictive model for sample 112, in other embodiments there may be multiple predictive models for sample 112. For example, different predictive models may be used to determine the model parameters for different portions of sample 112 (such as different organs or different types of tissue) and, thus, for different voxels. Therefore, in some embodiments different predictive models may be used to provide T_1 and T_2 values in different types of tissue, such as the values summarized in Table 1.

Tissue	T_1 (s)	T_2 (ms)
Cerebrospinal Fluid	0.8 - 20	110 - 2000
White Matter	0.76 - 1.08	61-100
Gray Matter	1.09 - 2.15	61 – 109
Meninges	0.5 - 2.2	50 – 165
Muscle	0.95 - 1.82	20 – 67
Adipose	0.2 - 0.75	53 – 94

Table 1

[061] Moreover, while system 100 is illustrated as having particular components, in other embodiments system 100 may have fewer or more components, two or more components may be combined into a single component, and/or positions of one or more components may be changed.

[062] We now embodiments of a method. FIG. 2 presents a flow diagram illustrating an example of a method 200 for determining model parameters associated with a sample. This method may be performed by a system (such as system 100 in FIG. 1), or one or more components in a system (such as source 110, measurement device 114 and/or computer 116).

[063] During operation, a source in the system may apply, to the sample, an excitation (operation 210), where the excitation has at least a wavelength and an intensity or a flux. For example, the excitation may include one of: electromagnetic radiation, a radio-frequency wave, a particle beam, a sound wave, a magnetic field, and/or an electric field. Therefore, the excitation may include at least one of: an electromagnetic beam in an x-ray band of wavelengths, a neutron beam, an electron beam, an electromagnetic beam in an optical band of wavelengths, an electromagnetic beam in an infrared band of wavelengths, a sound wave in an ultrasound band of wavelengths, a proton beam, an electric field associated with an impedance measurement device, a radio-frequency wave associated with a magnetic-resonance apparatus, and/or a magnetic field associated with a susceptibility measurement device.

[064] Then, a measurement device in the system may measure a response (operation 212) associated with the sample to the excitation. For example, the measurement device may include at least one of: an x-ray detector, a neutron detector, an electron detector, an optical detector, an infrared detector, an ultrasound detector, a proton detector, the magnetic-resonance apparatus, the impedance measurement device and/or the susceptibility measurement device. Note that the measured response may include a time-domain response of the sample and may be other than or different from an image.

[065] Moreover, the system may compute, using the measured response and information specifying the excitation as inputs to a predetermined predictive model, model parameters (operation 214) on a voxel-by-voxel basis in a forward model with multiple voxels that represent the sample. The forward model may simulate response physics occurring within the sample to a given excitation with a given wavelength and a given intensity or a given flux, that are selected from a range of measurement conditions that includes the excitation, the wavelength and the intensity or the flux, and at least a different wavelength and a at least a different intensity or a different flux. Furthermore, the forward model may be a function of

the excitation, the model parameters of the multiple voxels, and differential or phenomenological equations that approximates the response physics.

[066] Note that the predetermined predictive model may include a machine-learning model and/or a neural network. In some embodiments, the predetermined predictive model includes a personalized predictive model that corresponds to an individual.

[067] Next, the system may determine an accuracy of the model parameters (operation 216) by comparing at least the measured response and a calculated predicted value of the response using the forward model, the model parameters and the excitation.

[068] Additionally, when the accuracy exceeds a predefined value (operation 218), the system may provide the model parameters (operation 220) as, e.g., an output to a user, to another electronic device, to a display and/or to the memory.

[069] Thus, when the accuracy exceeds the predefined value (operation 218), the model parameters may be computed in a single pass without further iteration. Consequently, the model parameters having an accuracy exceeding the predefined value may be computed using fewer iterations with the predetermined predictive model than in the iterative approach without the predetermined predictive model.

[070] Alternatively, when the accuracy is less than the predefined value (operation 218), the system may: calculate, using information specifying the excitation and the accuracy as inputs to a second predetermined predictive model, a revised excitation (operation 222) that has at least a revised wavelength, a revised intensity or a revised flux; and repeat (operation 224) the applying, measuring, computing and determining operations with the revised excitation instead of the excitation. Note that the second predetermined predictive model may include a machine-learning model and/or a neural network.

[071] In some embodiments, the system optionally performs one or more optional additional or alternative operations. For example, when the accuracy is less than the predefined value (operation 218), the system may: add the excitation and the measured response to a training dataset; and determine, using the training dataset, a revised instance of the predictive model.

[072] Additionally, the system may classify or segment one or more anatomical structures in the sample using the model parameters and a third predictive model. For example, the

third predetermined predictive model may include a machine-learning model and/or a neural network.

[073] Moreover, the system may train the predictive model using a simulated dataset computed using the forward model, a range of model parameters and a range of excitations.

[074] In some embodiments of method 200, there may be additional or fewer operations. Furthermore, the order of the operations may be changed, and/or two or more operations may be combined into a single operation.

[075] FIG. 3 presents a drawing illustrating an example of communication among components in system 100 (FIG. 1). Notably, processor 310 in computer 116 may execute program instructions (P.I.) 312 stored in memory 314. When processor 310 executes program instructions 312, processor 310 may perform at least some of the operations in the computation technique.

[076] During the computation technique, processor 310 may provide instruction 318 to interface circuit (I.C.) 316. In response, interface circuit 316 may provide instruction 318 to source 110, e.g., in one or more packets or frames. Moreover, after receiving instructions 318, source 110 may apply, to the sample, an excitation 320.

[077] Then, processor 310 may provide instruction 322 to interface circuit 316. In response, interface circuit 316 may provide instruction 322 to measurement device 114, e.g., in one or more packets or frames. Furthermore, after receiving instructions 322, measurement device 114 may measure a response 324 associated with the sample to excitation 320. Next, measurement device 114 may provide measured response 324 to computer 116, e.g., in one or more packets or frames.

[078] After receiving measured response 324, interface circuit 316 may provide measured response 324 to processor 310. Then, using measured response 324 and information specifying excitation 320 as inputs to a predetermined predictive model, processor 310 may compute model parameters (M.P.) 326 on a voxel-by-voxel basis in a forward model with multiple voxels that represent the sample.

[079] Additionally, processor 310 may determine an accuracy 328 of the model parameters by comparing at least measured response 324 and a calculated predicted value of the response using the forward model, model parameters 326 and excitation 320. When accuracy 328 exceeds a predefined value, processor 310 may provide the model parameters 326 as, e.g., an

output to a user, to another electronic device (via interface circuit 316), to a display 330 and/or memory 314.

[080] Otherwise, when the accuracy is less than the predefined value, processor 310 may perform a remedial action 332. For example, processor 310 may: calculate, using information specifying excitation 320 and accuracy 328 as inputs to a second predetermined predictive model, a revised excitation; and repeat the applying, measuring, computing and determining operations with the revised excitation instead of excitation 320. Alternatively or additionally, processor 310 may: add excitation 320 and measured response 324 to a training dataset; and determine, using the training dataset, a revised instance of the predictive model.

[081] We now describe embodiments of predictive models. For example, a predictive model may include a machine-learning model, such as a supervised-learning model or an unsupervised learning technique (such as clustering). In some embodiments, a machine-learning model may include: a support vector machine, a classification and regression tree, logistic regression, LASSO, linear regression, nonlinear regression, pattern recognition, a Bayesian technique, and/or another (linear or nonlinear) supervised-learning technique.

[082] FIG. 4 presents a drawing illustrating an example of a machine-learning model 400. In this machine-learning model, a weighted (using weights 408) linear or nonlinear combination 416 of measurements 410, one or more corresponding excitations 412 and one or more errors 414 between the one or more measurements 410 and one or more predicted responses determined using a forward model, a current instance of the model parameters of voxels in the forward model, and the one or more excitations 412 is used to compute a revised instance of model parameters 418. Thus, in some embodiments, predictive model 400 is used in conjunction with forward model to iteratively modify instances of the model parameters until an accuracy of the predicted response is less than a predefined value (i.e., a convergence criterion is achieved). However, in some embodiments, a machine-learning model may be used to determine the model parameters in one pass, i.e., in an open-loop manner.

[083] Alternatively or additionally, a predictive model may include a neural network. Neural networks are generalized function approximators. For example, techniques such as deep learning typically use previous examples as inputs. In general, it is not possible for these machine-learning models to determine the actual function they are trying to

approximate because there is no reference point for them to use to estimate the error in their predictions. In particular, it can be difficult for a neural network to make predictions based on an input that is very different from the examples it was trained on. In this regard, a neural network can be thought of as a lossy compute compression engine.

[084] However, by training a neural network using a wide variety of excitations, measured responses and corresponding model parameters, the neural network can provide the model parameters (or initial estimates of the model parameters) for a forward model that simulates the physics of a response of a sample to an excitation. Because neural networks are effective approximations/compressions, they may execute faster on the same inputs with less computational power required. Moreover, because the functions are known in the forward model, the responses can be computed and the accuracy of the predictions can be assessed (as opposed to using an approximation). Therefore, the computation technique can be used to determine when its predictions are unreliable. In particular, as discussed previously for FIG. 4, a neural network may be used in conjunction with forward model to iteratively modify instances of the model parameters until an accuracy of the predicted response is less than a predefined value (i.e., a convergence criterion is achieved). In some embodiments, however, a neural network may be used to determine the model parameters in one pass, i.e., in an open-loop manner.

[085] FIG. 5 presents a drawing illustrating an example of a neural network 500. This neural network may be implemented using a convolutional neural network or a recurrent neural network. For example, neural network 500 may include a network architecture 512 that includes: an initial convolutional layer 514 that provides filtering of inputs 510 (such as one or more measurements and a difference or an error between the one or more measurements and one or more predicted responses determined using a forward model, a current instance of model parameters and an excitation); an additional convolutional layer(s) 516 that apply weights; and an output layer 518 (such as a rectified linear layer) that performs selection (*e.g.*, selecting a revised instance of the model parameters). Note that the details with the different layers in neural network 500, as well as their interconnections, may define network architecture 512 (such as a directed acyclic graph). These details may be specified by the instructions for neural network 500. In some embodiments, neural network 500 is reformulated as a series of matrix multiplication operations. Neural network 500 may be able

to handle the real-world variance in 1 million inputs or more. Note that neural network 500 may be trained using a deep-learning technique or a GAN. In some embodiments of machine-learning model 400 (FIG. 4) and/or neural network 500, a current instance of the model parameters is used as an input.

[086] In some embodiments, a large convolutional neural network may include 60 M parameters and 650,000 neurons. The convolutional neural network may include eight learned layers with weights, including five convolutional layers and three fully connected layers with a final 1000-way softmax or normalized exponential function that produces a distribution over the 1000 class labels for different possible model parameters. Some of the convolution layers may be followed by max-pooling layers. In order to make training faster, the convolutional neural network may use non-saturating neurons (such as a local response normalization) and an efficient dual parallelized GPU implementation of the convolution operation. In addition, in order to reduce overfitting in the fully-connected layers, a regularization technique (which is sometimes referred to as ‘dropout’) may be used. In dropout, the predictions of different models are efficiently combined to reduce test errors. In particular, the output of each hidden neuron is set to zero with a probability of 0.5. The neurons that are ‘dropped out’ in this way do not contribute to the forward pass and do not participate in backpropagation. Note that the convolutional neural network may maximize the multinomial logistic regression objective, which may be equivalent to maximizing the average across training cases of the log-probability of the correct label under the prediction distribution.

[087] In some embodiments, the kernels of the second, fourth, and fifth convolutional layers are coupled to those kernel maps in the previous layer that reside on the same GPU. The kernels of the third convolutional layer may be coupled to all kernel maps in the second layer. Moreover, the neurons in the fully connected layers may be coupled to all neurons in the previous layer. Furthermore, response-normalization layers may follow the first and second convolutional layers, and max-pooling layers may follow both response-normalization layers as well as the fifth convolutional layer. A nonlinear model of neurons, such as Rectified Linear Units, may be applied to the output of every convolutional and fully-connected layer.

[088] In some embodiments, the first convolutional layer filters a $224 \times 224 \times 3$ input image with 96 kernels of size $11 \times 11 \times 3$ with a stride of four pixels (this is the distance between the

receptive field centers of neighboring neurons in a kernel map). Note that the second convolutional layer may take as input the (response-normalized and pooled) output of the first convolutional layer and may filter it with 256 kernels of size $5 \times 5 \times 48$. Furthermore, the third, fourth, and fifth convolutional layers may be coupled to one another without any intervening pooling or normalization layers. The third convolutional layer may have 384 kernels of size $3 \times 3 \times 256$ coupled to the (normalized, pooled) outputs of the second convolutional layer. Additionally, the fourth convolutional layer may have 384 kernels of size $3 \times 3 \times 192$, and the fifth convolutional layer may have 256 kernels of size $3 \times 3 \times 192$. The fully-connected layers may have 4096 neurons each. Note that the numerical values in the preceding and the remaining discussion below are for purposes of illustration only, and different values may be used in other embodiments.

[089] In some embodiments, the convolutional neural network is implemented using at least two GPUs. One GPU may run some of the layer parts while the other runs the remaining layer parts, and the GPUs may communicate at certain layers. The input of the convolutional neural network may be 150,528-dimensional, and the number of neurons in the remaining layers in the convolutional neural network may be given by 253, 440–186, 624–64, 896–64, 896–43, and 264–4096–4096–1000.

[090] We now describe embodiments of a forward model. This forward model may be a 3D model of voxels in a portion of a sample (such as an individual), and may include model parameters in the Bloch equations for each of the voxels. In particular, with a quasi-static magnetic field B_0 along the z axis, the Bloch equations are

$$\begin{aligned}\frac{dM_x(t)}{dt} &= \gamma \cdot (\vec{M}(t) \otimes \vec{B}(t))_x - \frac{M_x(t)}{T_2}, \\ \frac{dM_y(t)}{dt} &= \gamma \cdot (\vec{M}(t) \otimes \vec{B}(t))_y - \frac{M_y(t)}{T_2}, \text{ and} \\ \frac{dM_z(t)}{dt} &= \gamma \cdot (\vec{M}(t) \otimes \vec{B}(t))_z - \frac{M_z(t) - M_0}{T_1},\end{aligned}$$

where γ is the gyromagnetic ratio, $\otimes$ denotes a vector cross product and

$\vec{B}(t) = (B_x(t), B_y(t), B_0 + \Delta B_z(t))$ is the magnetic field experienced by a type of nuclei in the sample. The model parameters in the Bloch equations may include T_1 , T_2 , a density of a type

of nuclei, diffusion, velocity/flow, temperature, magnetic susceptibility, etc. Note that there may be different model parameters for different types of nuclei for each of the voxels.

Moreover, note that the Bloch equations are a semi-classical, macroscopic approximation to the dynamic response of the magnetic moments of the type of nuclei in the sample to a time-varying magnetic field. For example, there may be 67 M cells in a 1 mm^3 voxel.

[091] In principle, the solution space for the model parameters in the Bloch equations for the sample may be underdetermined, i.e., there may be significantly more model parameters to be determined than there are observations with which to specify or constrain the model parameters. Therefore, when training a predictive model or determining the model parameters using the predictive model (such as using computations in a machine-learning model or in a layer in a neural network), the computation technique may leverage additional information to constrain or reduce the dimensionality of the problem. For example, an aspect of the anatomy of the sample may be determined using other imaging techniques, such as computed tomography, x-ray, ultrasound, etc. Moreover, regions that do not look like (i.e., that has very different measurements, e.g., different measured MR signals) than a targeted type of tissue (such as heart tissue) may be excluded from the forward model (such as by setting the model parameters to zero in these regions). In this way, e.g., regions that consist of air may be excluded. Other constraints in the forward model may include: thermodynamic constraints on heat flow (from hot to cold) for perfusion or MRT to quantify metabolism. In addition, the predictive model may be trained using measurements at different magnetic-field strengths B_0 (which may provide similar information to pseudorandom pulse sequences) using different pulse sequences and/or different MR techniques, which may reduce the ratio of model parameters to observations, thereby simplifying the training of the predictive model.

[092] Alternatively or additionally, tissue that deviates significantly from a predicted or simulated response (such as predicted MR signals) based on previous MR measurements or scans (e.g., anomalies or changes) may become the focus of the forward model, such as by using a contour map (e.g., a cubic spline) to bound the regions (or specify a boundary of the regions) where there are significant differences. In some embodiments, when training the predictive model or determining the model parameters using the predictive model (such as using computations in a machine-learning model or in a layer in a neural network), the difference or error between measurements and simulated or predicted responses may be

represented using one or more level-set functions, and the boundaries of regions with errors exceeding a threshold value may be determined based on the intersection of a plane corresponding to the threshold value and the one or more level-set functions.

[093] In some embodiments, a layer in a neural network may compute first and second derivatives along a surface(s) of model-parameter solutions in the sample. (In order to facilitate calculation of a derivative, the model parameters may be represented using one or more level-set functions.) A set of voxels along the line where the first derivative is zero may be identified. This set of voxels may be fit using a cubic spline with a minimum error between the voxel positions and the cubic spline. This fitting operation may be repeated at all the boundaries in the model-parameter-solution space. Moreover, the largest continuous surface within the boundary defined by the cubic splines may be determined and the model-parameter-solution calculation may be repeated to determine a new continuous surface that is within the previous continuous surface. This generalized framework may minimize the error across intra-voxel volumes, thereby improving the agreement between the measurements and the simulated or predicted responses based on the forward model.

[094] For example, a neural network may solve the inverse problem using a Jacobian matrix of the model parameters for the voxels in the forward model and Newton's method to modify the model parameters for the voxels in successive layers based on how perturbations in the model parameters affect the difference or error between the measurements and the predicted responses.

[095] In some embodiments, if a portion of the sample included one voxel, there may be 4-10 model parameters (which specify a forward model) that need to be determined for a particular type of tissue. If the voxel includes M types of tissue, there may be $4M$ - $10M$ the model parameters that need to be determined for the particular type of tissue. As the number of voxels increases, this can appear to be a daunting problem.

[096] However, because different types of nuclei have different Larmor frequencies, the spatial distribution of the types of nuclei and their local concentrations may be determined from the measurements. Then, a predefined anatomical template for the human body (or a portion of the human body), with associated initial model parameters for the forward model, may be scaled to match the spatial distribution of the types of nuclei and their local concentrations. For example, predetermined or predefined ranges for the model parameters

in different types of tissue may be used to determine for the initial model parameters. In some embodiments, the initial model parameters are based on model parameters associated with previous measurements or MR scans.

[097] Next, a look-up table with simulated or predicted responses (generated using one or more forward models) as a function of associated model parameters and excitations may be used to modify the initial model parameters or to compute model parameters for voxels in the sample. For example, simulated or predicted responses that are similar to measurements may be identified, and the differences or errors between these simulated or predicted responses and the measurements may be used to guide interpolation between the model parameters in the look-up table.

[098] In some embodiments, for a type of tissue (such as a particular organ), the model parameters determined using different layers in a neural network may be iteratively refined as the size of the voxels is progressively decreased (and, thus, the number of voxels is increased) in the different layers. This analysis may be driven by the error between the measurements and simulated or predicted responses using the forward model. Progressing through successive layers in a neural network, the focus may be on the residual regions with errors that are larger than a convergence or an accuracy criterion. For example, the model parameters for the forward model in a layer in a neural network may be based on measurements at one magnetic-field strength and then the error may be determined based on the predicted response of the forward model at another magnetic-field strength. Furthermore, note that initially the predictive model or the forward model may assume that there is no contribution or interaction between different voxels. However, as the error and the voxel size are reduced, such contributions and/or interactions may be included in subsequent layers in a neural network. In some embodiments, when there are multiple candidate model-parameter solutions (having similar errors) to the inverse problem for a layer in a neural network, at least some of these candidates may be kept for use in a subsequent layer (i.e., a unique model-parameter solution may not be identified at this point). Alternatively, if there is no unique parameter solution within a desired error range (such as less than 50, 25, 10, 5 or 1%), the best (least-error) model-parameter solution may be kept. In addition, when there is no model-parameter solution within the desired error range, the second predictive model may be used to modify the excitation and additional measurement(s) may be performed.

[099] Thus, the inverse problem of determining the model parameters based on measurements may be ‘solved’ using a predictive model that provides model parameters that minimize the error or difference between the measurements and simulated or predicted responses that are generated based on the forward model, the model parameters and an excitation. In some embodiments, the inverse problem is solved using one or more analysis techniques, including: a least-squares technique, a convex quadratic minimization technique, a steepest descents technique, a quasi-Newton technique, a simplex technique, a Levenberg–Marquardt technique, simulated annealing, a genetic technique, a graph-based technique, another optimization technique and/or Kalman filtering (or linear quadratic estimation).

[0100] Note that the training of a predictive model may use dynamic programming. In particular, the training problem may be divided up and performed by multiple computers in parallel, e.g., in a cloud-based computing system. For example, a particular thread may attempt to solve the inverse problem for particular measurement conditions. Multiple potential model-parameter solutions generated by the computers (or processors) may be combined (e.g., using linear superposition) to determine an error metric that is minimized using the one or more analysis techniques.

[0101] Moreover, as described previously, the inverse problem may be solved iteratively by a predictive model (such as machine-learning model or a neural network) by first attempting to find suitable model parameters (e.g., model parameters that minimize the error between measurements and simulated or predicted responses) for the forward model using a coarse voxel size and then progressively finding suitable parameters with smaller voxel sizes in subsequent layers or stages of the calculation. Note that the final voxel size used in this iterative procedure (or a suitable range of voxel sizes, because the voxel size may not be fixed in some embodiments) may be determined based on the gyromagnetic ratio of a type of nuclei being scanned. Furthermore, the voxel size or locations may also be chosen so that a voxel is evenly portioned into a set of subvoxels, or so that there is certain amount of overlap with previous voxel sizes to effectively ‘oversample’ the overlapping region and potentially further localize where an MR signal originates. This last technique may be akin to shifting the entire gradient system in one or more dimensions by a distance dx that is less than a characteristic length of the voxels (such as a length, a width or a height of the voxels). In some embodiments, the voxel size in the predictive model or the forward model is smaller

than that used in the measurements (i.e., the predictive model or the forward model may use a super-resolution technique). For example, there may be 512x512 voxels or 1024x1024 voxels at a magnetic-field strength of 3T. Note that the voxel size may be less than 0.25^3 mm³.

[0102] We now describe embodiments of a technique for segmenting different types of tissue, which may be used by the third predictive model (such as a neural network). Define a dictionary D_{mr} of measured time-sampled MR trajectories (or vectors) in a multi-dimensional parameter space for different types of tissue d_j (for $j=1$ to n) such that a measured MR signal y_{obv} for a voxel can be expressed as

$$y_{obv} = \sum_{j=1}^n \alpha_j \cdot d_j + \varepsilon,$$

where α_j are normalized weights (i.e., $\sum_{j=1}^n \alpha_j = 1$) and ε is an error (i.e., $\varepsilon = (y_j, \alpha_j)$, for j

$=1$ to n). This may define an intra-voxel linear equation problem. A generalized inter-voxel problem may model a set of voxels (such as a cube with 27 voxels) as a graph G . Note that each voxel in the set may have 26 edges to eight adjacent voxels. A parameter solution to the inverse problem may be defined as one that minimizes the error.

[0103] Consider the case of two adjacent voxels u and v . The intra-voxel linear equations U_y and V_y need to be solved at both u and v . There are several possible outcomes. First, U_y and V_y may have unique model-parameter solutions (where a ‘unique model-parameter solution’ may be a best fit to an existing forward model, i.e., with an error or difference vector that is less than a convergence or an accuracy criterion) and the analysis may be finished.

Alternatively, U_y may have a unique model-parameter solution but not V_y . It may be possible that the model-parameter solution for U_y imposes a constraint on V_y such that V_y has a single model-parameter solution, in which case the analysis may be finished. However, neither U_y and V_y may have unique model-parameter solutions, in which case combining the systems of equations (i.e., effectively increasing the voxel size) may yield a unique model-parameter solution. Moreover, neither U_y and V_y may have any model-parameter solutions, in which case the intra-voxel problem cannot be solved without further constraints.

[0104] In the last case, it may be possible to look at an adjacent voxel w , i.e., series voxels u , v and w , with the corresponding intra-voxel linear equations U_y , V_y and W_y need to be solved

at u, v and w. Note that the intra-voxel linear equations V_y and W_y reduce to the previous case. When the intra-voxel linear equations do not reduce to the previous case, this paring operation can be applied recursively until it does and then the intra-voxel linear equations can be solved as described previously.

[0105] In general, this analysis technique may be isomorphic to the problem of fitting a 3D surface (or volume) to minimize the error. One challenge in this regard is that it assumes that all adjacent volumes have an equal effect on the model-parameter solution α_j that minimizes the error.

[0106] The minimization of the error may initially assume that there is no inter-voxel contribution (i.e., that the voxels are independent). Subsequently, inter-voxel contributions may be included. In particular, considering adjacent voxel volumes, there are two distinct classes. Volumes that share a surface and volumes that only share a 1D edge. The minimization function can be improved by weighting the error contribution at voxel u at the center of the relative co-ordinate system. If the effect on the error is proportional to r^{-2} (where r is the distance between center points of voxels) and assuming 1 mm isotropic voxels in the weightings, the minimization or fitting problem with inter-voxel contributions can be expressed as

$$\min(\text{error}(y(0,0,0), \alpha(0,0,0)) + \frac{1}{(1)^2} \sum_{k=1}^m \text{error}(y_k, \alpha_k) + \frac{1}{(\sqrt{2})^2} \sum_{l=1}^p \text{error}(y_l, \alpha_l),$$

where the summation over k is for adjacent voxels sharing a common surface (i.e., (-1,0,0), (1,0,0), (0,-1,0), (0,1,0), (0,0,-1) and (0,0,1)) and the summation over l is for a remainder of adjacent voxels sharing a common edge. The assumption in the analysis is that the most difficult place to fit or determine model-parameter solutions is at discontinuities or interfaces between different tissues. Consequently, during the computation technique, analysis engine 122 (FIG. 1) may solve these locations first and then may solve the remaining locations.

[0107] Alternatively, because the magnetic contribution from neighboring voxels is proportional to r^2 , given a sphere of radius R from the center of a primary or central voxel in the minimization problem, surrounding voxels may be weighted based on the how much the sphere expands into the volume of the adjacent voxels (and, thus, based on how strong their inter-voxel contribution is estimated to be). For example, there may be three different weights that need to be assigned, including: a weight for voxels that share a 2D surface, a

weight for voxels that share a 1D line, and a weight for voxels that share a 0D point. Because there may not be a uniform tissue distribution within each voxel, the weights may be dynamically adjusted to model different kinds of distributions inside each voxel in order find the distributions that minimize the error. This may provide the ability to identify multiple MR signatures within a single voxel for different types of tissue. Note that, as computational power increases, the accuracy of the third predictive model may increase and the analysis technique used to solve the minimization problem (and, thus, the inverse problem) may be modified.

[0108] Thus, in embodiments where the forward model of a voxel depends on the forward models of surrounding or neighboring voxels, the forward model of a voxel may be computed using 2nd or Nth-order effects. For example, if there are N 1st-order forward models (where N is an integer), there may be as many as $N!/(N-27)!$ 2nd-order forward models (if all the voxels interact with each other). In some embodiments, locality is used to simplify the inverse problem. In this way, a forward model may be generated by incorporating how the forward models in adjacent voxels effect the forward model in a primary (central) or 1st-order voxel.

[0109] In some embodiments, a dithering technique is used to overcome the arbitrary locations of the voxels relative to the distribution of types of tissue in the body. In particular, there may be two or more types of tissue in a voxel because of the arbitrary voxel placement or the current voxel size. This may significantly change the forward model parameters for this voxel. This may suggest that there is more than one forward model needed for the voxel. In order to confirm this, the voxels may be displaced by a distance dx (which is a fraction of the voxel length, width or height) and the forward model parameters may be determined again (e.g., using the predictive model). In the processes, the tissue distribution may be determined. Consequently, this approach may effectively increase the spatial resolution in the analysis without changing the voxel size.

[0110] FIG. 6 presents a drawing illustrating an example of classification or segmentation of one or more anatomical structures 600. Notably, FIG. 6 illustrates identifying or segmenting an organ 610 based at least in part on discontinuous changes in T₁ and T₂ at voxel boundaries.

[0111] While the preceding discussion illustrated the computation technique using MR techniques, this approach may be generalized to a measurement system that is able to

physically model and measure a sample in real-time using a wide variety of characterization techniques. In general, the computation technique can use a combination of mechanical and/or electromagnetic waves to ‘perturb’ or ‘excite’ the volume being scanned in order to evaluate the correctness of a prediction in terms of how the volume responds to the perturbation. This also includes the ability for the system to simulate itself and any part of the environment in which the system is located that could affect the correctness or accuracy of the forward model the system is trying to generate to describe the volume being scanned or measured.

[0112] Note that the different characterization techniques may provide tensor-field mapping and the ability to detect anomalies in tensor fields. These maps can be images or quantitative tensor field maps, and each of the characterization techniques may provide a visualization of a different type of tensor field map captured with different type of measurements. By looking at or considering two or more of these maps, of the system may have access to orthogonal information.

[0113] Thus, the system may provide a way to capture, in real-time or near real-time, higher-order or hyper-dimensional pseudo or hybrid tensors or matrices at each voxel in 3D space. Using electromagnetic and/or mechanical perturbations or excitations, the system may use different characterization techniques to measure disturbances and responses, and then to simulate the responses.

[0114] The result of this characterization may be a $(4+N)D$ (three spatial dimensions, one time dimension, and up to N measurement dimensions at each point in space) quantitative model of the volume being scanned. Note that the $(4+N)D$ quantitative model may be projected onto an arbitrary subset of the full $(4+N)D$ space, including 2D or 3D images.

[0115] In some embodiments, the use of multidimensional data and models provides enhanced diagnostic accuracy (i.e., a lower false-positive rate) relative to conventional MRI approaches, even if a larger voxel size is used. Thus, the computation technique may allow improved diagnostic accuracy with a larger voxel size (or a weaker external magnetic field) than would be needed in conventional MRI. However, as noted previously, the computation technique may be used with a wide variety of measurement techniques separately from or in addition to MRI.

[0116] We now further describe an electronic device that performs at least some of the operations in the computation technique. FIG. 7 presents a block diagram illustrating an electronic device 700 in system 100 (FIG. 1), such as computer 116 (FIG. 1) or another of the computer-controlled components in system 100, such as source 110 or measurement device 114 (FIG. 1). This electronic device includes a processing subsystem 710, memory subsystem 712, and networking subsystem 714. Processing subsystem 710 may include one or more devices configured to perform computational operations and to control components in system 100 (FIG. 1). For example, processing subsystem 710 may include one or more microprocessors or central processing units (CPUs), one or more graphics processing units (GPUs), application-specific integrated circuits (ASICs), microcontrollers, programmable-logic devices (such as a field programmable logic array or FPGA), and/or one or more digital signal processors (DSPs).

[0117] Memory subsystem 712 may include one or more devices for storing data and/or instructions for processing subsystem 710 and networking subsystem 714. For example, memory subsystem 712 may include dynamic random access memory (DRAM), static random access memory (SRAM), and/or other types of memory. In some embodiments, instructions for processing subsystem 710 in memory subsystem 712 include one or more program modules or sets of instructions (such as program instructions 724), which may be executed in an operating environment (such as operating system 722) by processing subsystem 710. Note that the one or more computer programs may constitute a computer-program mechanism or a program module (i.e., software). Moreover, instructions in the various modules in memory subsystem 712 may be implemented in: a high-level procedural language, an object-oriented programming language, and/or in an assembly or machine language. Furthermore, the programming language may be compiled or interpreted, e.g., configurable or configured (which may be used interchangeably in this discussion), to be executed by processing subsystem 710.

[0118] In addition, memory subsystem 712 may include mechanisms for controlling access to the memory. In some embodiments, memory subsystem 712 includes a memory hierarchy that comprises one or more caches coupled to a memory in electronic device 700. In some of these embodiments, one or more of the caches is located in processing subsystem 710.

[0119] In some embodiments, memory subsystem 712 is coupled to one or more high-capacity mass-storage devices (not shown). For example, memory subsystem 712 may be coupled to a magnetic or optical drive, a solid-state drive, or another type of mass-storage device. In these embodiments, memory subsystem 712 may be used by electronic device 700 as fast-access storage for often-used data, while the mass-storage device is used to store less frequently used data.

[0120] In some embodiments, memory subsystem 712 includes a remotely located archive device. This archive device can be a high-capacity network attached mass-storage device, such as: network attached storage (NAS), an external hard drive, a storage server, a cluster of servers, a cloud-storage provider, a cloud-computing provider, a magnetic-tape backup system, a medical records archive service, and/or another type of archive device. Moreover, processing subsystem 710 may interact with the archive device via an application programming interface to store and/or access information from the archive device. Note that memory subsystem 712 and/or electronic device 700 may be compliant with the Health Insurance Portability and Accountability Act.

[0121] An example of the data stored (locally and/or remotely) in memory subsystem 712 is shown in FIG. 8, which presents a drawing illustrating an example of a data structure 800 that is used by electronic device 700 (FIG. 7). This data structure may include: an identifier 810-1 of sample 808-1 (such as an individual), metadata 812 (such as age, gender, biopsy results and diagnosis if one has already been made, other sample information, demographic information, family history, etc.), timestamps 814 when data was acquired, received measurements 816 (such as MR signals and, more generally, raw data), excitation and measurement conditions 818 (such as an external magnetic field, an optional gradient, an RF pulse sequence, an MR apparatus, a location, machine-specific characteristics such as magnetic-field inhomogeneity, RF noise and one or more other system imperfections, signal-processing techniques, registration information, synchronization information such between measurements and a heartbeat or breathing pattern of an individual, etc.), and/or determined model parameters 820 (including voxel sizes, speed, resonant frequency or a type of nuclei, T_1 and T_2 relaxation times, segmentation information, classification information, etc.), environmental conditions 822 (such as the temperature, humidity and/or barometric pressure in the room or the chamber in which sample 808-1 was measured), forward model 824, one

or more additional measurements 826 of physical properties of sample 808-1 (such as weight, dimensions, images, etc.), optional detected anomalies 828 (which may include particular voxel(es) associated with the one or more of detected anomalies 828), and/or optional classifications 830 of the one or more detected anomalies 828. Note that data structure 800 may include multiple entries for different measurements.

[0122] In one embodiment, data in data structure 800 is encrypted using a block-chain or a similar cryptographic hash technique to detect unauthorized modification or corruption of records. Moreover, the data can be anonymized prior to storage so that the identity of an individual associated with a sample is anonymous unless the individual gives permission or authorization to access or release the individual's identity.

[0123] Referring back to FIG. 7, networking subsystem 714 may include one or more devices configured to couple to and communicate on a wired, optical and/or wireless network (i.e., to perform network operations and, more generally, communication), including: control logic 716, an interface circuit 718, one or more antennas 720 and/or input/output (I/O) port 728. (While FIG. 7 includes one or more antennas 720, in some embodiments electronic device 700 includes one or more nodes 708, e.g., a pad or connector, which can be coupled to one or more antennas 720. Thus, electronic device 700 may or may not include one or more antennas 720.) For example, networking subsystem 714 can include a Bluetooth networking system (which can include Bluetooth Low Energy, BLE or Bluetooth LE), a cellular networking system (e.g., a 3G/4G network such as UMTS, LTE, etc.), a universal serial bus (USB) networking system, a networking system based on the standards described in IEEE 802.11 (e.g., a Wi-Fi networking system), an Ethernet networking system, and/or another networking system.

[0124] Moreover, networking subsystem 714 may include processors, controllers, radios/antennas, sockets/plugs, and/or other devices used for coupling to, communicating on, and handling data and events for each supported networking system. Note that mechanisms used for coupling to, communicating on, and handling data and events on the network for each network system are sometimes collectively referred to as a 'network interface' for network subsystem 714. Moreover, in some embodiments a 'network' between components in system 100 (FIG. 1) does not yet exist. Therefore, electronic device 700 may use the mechanisms in networking subsystem 714 for performing simple wireless communication

between the components, e.g., transmitting advertising or beacon frames and/or scanning for advertising frames transmitted by other components.

[0125] Within electronic device 700, processing subsystem 710, memory subsystem 712, networking subsystem 714 may be coupled using one or more interconnects, such as bus 726. These interconnects may include an electrical, optical, and/or electro-optical connection that the subsystems can use to communicate commands and data among one another. Although only one bus 726 is shown for clarity, different embodiments can include a different number or configuration of electrical, optical, and/or electro-optical connections among the subsystems.

[0126] Electronic device 700 may be (or can be) included in a wide variety of electronic devices. For example, electronic device 700 may be included in: a tablet computer, a smartphone, a smartwatch, a portable computing device, a wearable device, test equipment, a digital signal processor, a cluster of computing devices, a laptop computer, a desktop computer, a server, a subnotebook/netbook and/or another computing device.

[0127] Although specific components are used to describe electronic device 700, in alternative embodiments, different components and/or subsystems may be present in electronic device 700. For example, electronic device 700 may include one or more additional processing subsystems, memory subsystems, and/or networking subsystems. Additionally, one or more of the subsystems may not be present in electronic device 700. Moreover, in some embodiments, electronic device 700 may include one or more additional subsystems that are not shown in FIG. 7.

[0128] Although separate subsystems are shown in FIG. 7, in some embodiments, some or all of a given subsystem or component can be integrated into one or more of the other subsystems or components in electronic device 700. For example, in some embodiments program instructions 724 are included in operating system 722. In some embodiments, a component in a given subsystem is included in a different subsystem. Furthermore, in some embodiments electronic device 700 is located at a single geographic location or is distributed over multiple different geographic locations.

[0129] Moreover, the circuits and components in electronic device 700 may be implemented using any combination of analog and/or digital circuitry, including: bipolar, PMOS and/or NMOS gates or transistors. Furthermore, signals in these embodiments may include digital

signals that have approximately discrete values and/or analog signals that have continuous values. Additionally, components and circuits may be single-ended or differential, and power supplies may be unipolar or bipolar.

[0130] An integrated circuit may implement some or all of the functionality of networking subsystem 714 (such as a radio) and, more generally, some or all of the functionality of electronic device 700. Moreover, the integrated circuit may include hardware and/or software mechanisms that are used for transmitting wireless signals from electronic device 700 and receiving signals at electronic device 700 from other components in system 100 (FIG. 1) and/or from electronic devices outside of system 100 (FIG. 1). Aside from the mechanisms herein described, radios are generally known in the art and hence are not described in detail. In general, networking subsystem 714 and/or the integrated circuit can include any number of radios. Note that the radios in multiple-radio embodiments function in a similar way to the radios described in single-radio embodiments.

[0131] While some of the operations in the preceding embodiments were implemented in hardware or software, in general the operations in the preceding embodiments can be implemented in a wide variety of configurations and architectures. Therefore, some or all of the operations in the preceding embodiments may be performed in hardware, in software or both.

[0132] In addition, in some of the preceding embodiments there are fewer components, more components, a position of a component is changed and/or two or more components are combined.

[0133] While the preceding discussion illustrated the computation technique to solve a vector wave equation, in other embodiments the computation technique may be used to solve a scalar equation. For example, an acoustic wave equation may be solved in an arbitrary inhomogeneous media based on ultrasound measurements using a forward model. (Thus, in some embodiments the excitation may be mechanical.) Note that the acoustic coupling in ultrasound measurements can dependent on the operator (i.e., the ultrasound measurements may be pressure dependent). Nonetheless, a similar approach may be used to: improve ultrasound imaging, determine 3D structure, facilitate improved presentation, etc.

[0134] In the preceding description, we refer to ‘some embodiments.’ Note that ‘some embodiments’ describes a subset of all of the possible embodiments, but does not always

specify the same subset of embodiments. Moreover, note that numerical values in the preceding embodiments are illustrative examples of some embodiments. In other embodiments of the computation techniques, different numerical values may be used.

[0135] The foregoing description is intended to enable any person skilled in the art to make and use the disclosure, and is provided in the context of a particular application and its requirements. Moreover, the foregoing descriptions of embodiments of the present disclosure have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present disclosure to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present disclosure. Additionally, the discussion of the preceding embodiments is not intended to limit the present disclosure. Thus, the present disclosure is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

CLAIMS

1. A method for determining model parameters associated with a sample, comprising operations of:

5 applying, to the sample, an excitation using a source, in which the excitation has at least a wavelength and an intensity or a flux, and in which the excitation comprises one of: electromagnetic radiation, a radio-frequency wave, a particle beam, a sound wave, a magnetic field, or an electric field;

 measuring, using a measurement device, a response associated with the sample to the excitation;

10 computing, using the measured response and information specifying the excitation as inputs to a predetermined predictive model, the model parameters on a voxel-by-voxel basis in a forward model with multiple voxels that represent the sample, in which the forward model simulates response physics occurring within the sample to a given excitation with a given wavelength and a given intensity or a given flux, that are
15 selected from a range of measurement conditions that includes the excitation, the wavelength and the intensity or the flux, and at least a different wavelength and at least a different intensity or a different flux, and

 in which the forward model is a function of the excitation, the model parameters of the multiple voxels, and differential or phenomenological equations that approximates the
20 response physics;

 determining, using a computer, an accuracy of the model parameters by comparing at least the measured response and a calculated predicted value of the response using the forward model, the model parameters and the excitation;

25 providing, when the accuracy exceeds a predefined value, the model parameters as an output to a user, to another electronic device, to a display or to memory; and

 when the accuracy is less than the predefined value, modifying the model parameters, and repeating the computing and the determining.

2. A method according to Claim 1, in which the predetermined predictive model comprises one of: a machine-learning model, or a neural network.

3. A method according to Claim 1, in which the predetermined predictive model comprises a personalized predictive model that corresponds to an individual.
4. A method according to Claim 1, in which, when the accuracy exceeds the predefined value, the model parameters are computed in a single pass without further iteration.
- 5 5. A method according to Claim 1, in which the model parameters having the accuracy exceeding the predefined value are computed using fewer iterations with the predetermined predictive model than an iterative measurement and analysis technique without the predetermined predictive model.
6. A method according to Claim 1, in which, when the accuracy is less than the predefined value, the method comprises:
10 calculating, using information specifying the excitation and the accuracy as inputs to a second predetermined predictive model, a revised excitation that has at least a revised wavelength, a revised intensity or a revised flux; and
repeating the applying, measuring, computing and determining operations with the revised excitation instead of the excitation.
15
7. A method according to Claim 6, in which the second predetermined predictive model comprises one of: a machine-learning model, or a neural network.
8. A method according to Claim 1, in which, when the accuracy is less than the predefined value, the method comprises:
20 adding the excitation and the measured response to a training dataset; and
determining, using the training dataset, a revised instance of the predetermined predictive model.
9. A method according to Claim 1, in which the method comprises classifying or segmenting one or more anatomical structures in the sample using the model parameters and a
25 second predetermined predictive model.
10. A method according to Claim 9, in which the second predetermined predictive model comprises one of: a machine-learning model, or a neural network.

11. A method according to Claim 1, in which the method comprises training the predetermined predictive model using a simulated dataset computed using the forward model, a range of model parameters and a range of excitations.

12. A method according to Claim 1, in which the excitation is one of: an electromagnetic beam in an x-ray band of wavelengths, a neutron beam, an electron beam, an electromagnetic beam in an optical band of wavelengths, an electromagnetic beam in an infrared band of wavelengths, a sound wave in an ultrasound band of wavelengths, a proton beam, the electric field associated with an impedance measurement device, the radio-frequency wave associated with a magnetic-resonance apparatus, or the magnetic field associated with a susceptibility measurement device.

13. A method according to Claim 1, in which the measured response comprises a time-domain response of the sample and is other than an image.

14. A non-transitory computer-readable storage medium for use in conjunction with a computer, the computer-readable storage medium configured to store program instructions that, when executed by the computer, cause the computer to perform operations comprising:

applying, to a sample, an excitation using a source, in which the excitation has at least a wavelength and an intensity or a flux, and in which the excitation comprises one of: electromagnetic radiation, a radio-frequency wave, a particle beam, a sound wave, a magnetic field, or an electric field;

measuring, using a measurement device, a response associated with the sample to the excitation;

computing, using the measured response and information specifying the excitation as inputs to a predetermined predictive model, model parameters on a voxel-by-voxel basis in a forward model with multiple voxels that represent the sample, in which the forward model simulates response physics occurring within the sample to a given excitation with a given wavelength and a given intensity or a given flux, that are selected from a range of measurement conditions that includes the excitation, the wavelength and the intensity or the flux, and at least a different wavelength and at least a different intensity or a different flux, and

in which the forward model is a function of the excitation, the model parameters of the multiple voxels, and differential or phenomenological equations that approximates the response physics;

determining, using the computer, an accuracy of the model parameters by comparing at least the measured response and a calculated predicted value of the response using the forward model, the model parameters and the excitation;

providing, when the accuracy exceeds a predefined value, the model parameters as an output to a user, to another electronic device, to a display or to memory; and

when the accuracy is less than the predefined value, modifying the model parameters, and repeating the computing and the determining.

15. A non-transitory computer-readable storage medium according to Claim 14, in which the predetermined predictive model comprises one of: a machine-learning model, or a neural network.

16. A non-transitory computer-readable storage medium according to Claim 14, in which, when the accuracy exceeds the predefined value, the model parameters are computed in a single pass without further iteration.

17. A non-transitory computer-readable storage medium according to Claim 14, in which, when the accuracy is less than the predefined value, the operations comprise:

calculating, using information specifying the excitation and the accuracy as inputs to a second predetermined predictive model, a revised excitation that has at least a revised wavelength, a revised intensity or a revised flux; and

repeating the applying, measuring, computing and determining operations with the revised excitation instead of the excitation.

18. A non-transitory computer-readable storage medium according to Claim 17, in which the second predetermined predictive model comprises one of: a machine-learning model, or a neural network.

19. A system configured to determine model parameters, comprising:

a source configured to provide an excitation, in which the excitation comprises at least one of: an electromagnetic beam in an x-ray band of wavelengths, a neutron beam, an

electron beam, an electromagnetic beam in an optical band of wavelengths, an
electromagnetic beam in an infrared band of wavelengths, a sound wave in an
ultrasound band of wavelengths, a proton beam, an electric field associated with an
impedance measurement device, a radio-frequency wave associated with a magnetic-
5 resonance apparatus, or a magnetic field associated with a susceptibility measurement
device;

a measurement device configured to perform measurements, in which the measurement
device comprises at least one of: an x-ray detector, a neutron detector, an electron
detector, an optical detector, an infrared detector, an ultrasound detector, a proton
10 detector, the magnetic-resonance apparatus, the impedance measurement device or the
susceptibility measurement device;

a processor, coupled to the source, the measurement device and memory, configured to
execute program instructions; and

the memory, coupled to the processor, configured to store the program instructions that,
15 when executed by the processor, cause the system to perform operations comprising:

applying, to a sample, the excitation using the source, in which the excitation has
at least a wavelength and an intensity or a flux, and in which the excitation
comprises one of: electromagnetic radiation, a radio-frequency wave, a particle
beam, a sound wave, a magnetic field, or an electric field;

20 measuring, using a measurement device, a response associated with the sample to
the excitation;

computing, using the measured response and information specifying the excitation
as inputs to a predetermined predictive model, model parameters on a voxel-by-
voxel basis in a forward model with multiple voxels that represent the sample, in
25 which the forward model simulates response physics occurring within the sample
to a given excitation with a given wavelength and a given intensity or a given flux,
that are selected from a range of measurement conditions that includes the
excitation, the wavelength and the intensity or the flux, and at least a different
wavelength and at least a different intensity or a different flux, and

in which the forward model is a function of the excitation, the model parameters of the multiple voxels, and differential or phenomenological equations that approximates the response physics;

5 determining, using the processor, an accuracy of the model parameters by comparing at least the measured response and a calculated predicted value of the response using the forward model, the model parameters and the excitation;

providing, when the accuracy exceeds a predefined value, the model parameters as an output to a user, to another electronic device, to a display or to the memory; and

10 when the accuracy is less than the predefined value, modifying the model parameters, and repeating the computing and the determining.

20. A system according to Claim 19, in which the predetermined predictive model comprises one of: a machine-learning model, or a neural network.