

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2020206584 B2**

(54) Title
Systems and methods for platform agnostic whole body image segmentation

(51) International Patent Classification(s)
G06T 7/11 (2017.01)

(21) Application No: **2020206584** (22) Date of Filing: **2020.01.06**

(87) WIPO No: **WO20/144134**

(30) Priority Data

(31) Number	(32) Date	(33) Country
62/789,155	2019.01.07	US
62/837,941	2019.04.24	US
62/863,608	2019.06.19	US
62/870,210	2019.07.03	US
62/907,158	2019.09.27	US
62/934,305	2019.11.12	US

(43) Publication Date: **2020.07.16**

(44) Accepted Journal Date: **2024.11.07**

(71) Applicant(s)
EXINI Diagnostics AB

(72) Inventor(s)
RICHTER, Jens Filip Andreas;JOHNSSON, Kerstin Elsa Maria;GJERTSSON, Erik Konrad;ANAND, Aseem Undvall

(74) Agent / Attorney
Pizzey's Patent and Trade Mark Attorneys Pty Ltd, PO Box 291, WODEN, ACT, 2606, AU

(56) Related Art
GJERTSSON, K. "Segmentation in Skeletal Scintigraphy Images using Convolutional Neural Networks", MASTER'S THESES IN MATHEMATICAL SCIENCES, 22 June 2017

**PEIRUI BAI ET AL: "Body region localization in whole-body low-dose CT images of PET/CT scans using virtual landmarks", MEDICAL PHYSICS WILEY USA, vol. 46, no. 3, 4 January 2019 (2019-01-04), pages 1286 - 1299, XP055672103, ISSN: 0094-2405
US 8855387 B2**

(51) International Patent Classification:
G06T 7/11 (2017.01)

(21) International Application Number:
PCT/EP2020/050132

(22) International Filing Date:
06 January 2020 (06.01.2020)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

62/789,155	07 January 2019 (07.01.2019)	US
62/837,941	24 April 2019 (24.04.2019)	US
62/863,608	19 June 2019 (19.06.2019)	US
62/870,210	03 July 2019 (03.07.2019)	US
62/907,158	27 September 2019 (27.09.2019)	US
62/934,305	12 November 2019 (12.11.2019)	US

(71) Applicant: **EXINI DIAGNOSTICS AB** [SE/SE]; Ideon Science Park, Scheelevägen 27, Gateway, SE-223 70 Lund (SE).

(72) Inventors: **RICHTER, Jens Filip Andreas**; c/o EXINI Diagnostic AB, Ideon Science Park, Scheelevägen 27, Gateway, SE-223 70 Lund (SE). **JOHNSSON, Kerstin Elsa Maria**; c/o EXINI Diagnostics AB, Ideon Science Park, Scheelevägen 27, Gateway, SE-223 70 Lund (SE). **GJERTSSON, Erik Konrad**; c/o EXINI Diagnostics AB, Ideon Science Park, Scheelevägen 27, Gateway, SE-223 70 Lund (SE). **ANAND, Aseem Undvall**; c/o EXINI Diagnostics AB, Ideon Science Park, Scheelevägen 27, Gateway, SE-223 70 Lund (SE).

(74) Agent: **GRAHAM WATT & CO LLP**; St Botolph's House, 7-9 St Botolph's Road, Sevenoaks Kent TN13 3AJ (GB).

(54) Title: SYSTEMS AND METHODS FOR PLATFORM AGNOSTIC WHOLE BODY IMAGE SEGMENTATION




FIG. 4

(57) Abstract: Presented herein are systems and methods that provide for automated analysis of three-dimensional (3D) medical images of a subject in order to automatically identify specific 3D volumes within the 3D images that correspond to specific anatomical regions (e.g., organs and/or tissue). Notably, the image analysis approaches described herein are not limited to a single particular organ or portion of the body. Instead, they are robust and widely applicable, providing for consistent, efficient, and accurate detection of anatomical regions, including soft tissue organs, in the entire body. In certain embodiments, the accurate identification of one or more such volumes is used to automatically determine quantitative metrics that represent uptake of radiopharmaceuticals in particular organs and/or tissue regions. These uptake metrics can be used to assess disease state in a subject, determine a prognosis for a subject, and/or determine efficacy of a treatment modality.

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

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

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

SYSTEMS AND METHODS FOR PLATFORM AGNOSTIC WHOLE BODY IMAGE SEGMENTATION**CROSS-REFERENCE TO RELATED APPLICATIONS**

[0001] This application claims priority to and benefit of U.S. Provisional Application No. 62/789,155, filed January 07, 2019, U.S. Provisional Application No. 62/837,941, filed April 24, 2019, U.S. Provisional Application No. 62/863,608, filed June 19, 2019, U.S. Provisional Application No. 62/870,210, filed July 03, 2019, U.S. Provisional Application No. 62/907,158, filed September 27, 2019, U.S. Provisional Application No. 62/934,305, filed November 12, 2019, the content of each of which are hereby incorporated by reference in their entirety.

FIELD OF THE INVENTION

[0002] This invention relates generally to methods, systems, and architectures for automated analysis of medical images. More particularly, in certain embodiments, the invention relates to automated identification of one or more particular regions of interest (e.g., corresponding to specific organs or tissue) within images of a subject.

BACKGROUND OF THE INVENTION

[0003] Targeted image analysis involves the use of radiolabeled small molecules that bind to specific receptors, enzymes and proteins in the body that are altered during the evolution of disease. After administration to a patient, these molecules circulate in the blood until they find their intended target. The bound radiopharmaceutical remains at the site of disease, while the rest of the agent clears from the body. The radioactive portion of the molecule serves as a beacon so that an image may be obtained depicting the disease location and concentration using commonly available nuclear medicine cameras, known as single-photon emission computerized tomography (SPECT) or positron emission tomography (PET) cameras, found in most hospitals throughout the world. Physicians can then use this information to determine the presence and the extent of disease in a patient. The physician can use this information to provide a recommended course of treatment to the patient and to track the progression of disease.

[0004] There are a variety of software-based analytical techniques available for analysis and enhancement of PET and SPECT images that can be used by a radiologist or physician. There are also a number of radiopharmaceuticals available for imaging particular kinds of cancer. Imaging agents used in the art include, among others include, without limitation ^{18}F -NaF, ^{11}C -Choline, 2-deoxy-2 [^{18}F] fluoro-d-glucose (FDG), and the like. For example, the small molecule diagnostic 1404 targets the extracellular domain of prostate specific membrane antigen (PSMA), a protein amplified on the

surface of >95% of prostate cancer cells and a validated target for the detection of primary and metastatic prostate cancer. 1404 is labeled with technetium-99m, a gamma-emitter isotope that is widely available, relatively inexpensive, facilitates efficient preparation, and has spectrum characteristics attractive for nuclear medicine imaging applications.

[0005] Another example radiopharmaceutical is PyL™ (also known as [¹⁸F]DCFPyL, and 18F-PyL), which is a clinical-stage, fluorinated PSMA-targeted PET imaging agent for prostate cancer. A proof-of-concept study published in the April 2016 issue of the Journal of Molecular Imaging and Biology demonstrated that PET imaging with PyL™ showed high levels of PyL™ uptake in sites of putative metastatic disease and primary tumors, suggesting the potential for high sensitivity and specificity in detecting prostate cancer.

[0006] An oncologist may use images from a targeted PET or SPECT study of a patient as input in her assessment of whether the patient has a particular disease, e.g., prostate cancer, what stage of the disease is evident, what the recommended course of treatment (if any) would be, whether surgical intervention is indicated, and likely prognosis. The oncologist may use a radiologist report in this assessment. A radiologist report is a technical evaluation of the PET or SPECT images prepared by a radiologist for a physician who requested the imaging study and includes, for example, the type of study performed, the clinical history, a comparison between images, the technique used to perform the study, the radiologist's observations and findings, as well as overall impressions and recommendations the radiologist may have based on the imaging study results. A signed radiologist report is sent to the physician ordering the study for the physician's review, followed by a discussion between the physician and patient about the results and recommendations for treatment.

[0007] Thus, the process involves having a radiologist perform an imaging study on the patient, analyzing the images obtained, creating a radiologist report, forwarding the report to the requesting physician, having the physician formulate an assessment and treatment recommendation, and having the physician communicate the results, recommendations, and risks to the patient. The process may also involve repeating the imaging study due to inconclusive results, or ordering further tests based on initial results.

[0008] If an imaging study shows that the patient has a particular disease or condition (e.g., cancer), the physician discusses various treatment options, including surgery, as well as risks of doing nothing or adopting a watchful waiting or active surveillance approach, rather than having surgery.

[0009] There are limitations associated with this process, both from the perspective of the physician and from the perspective of the patient. While the radiologist's report is certainly helpful, the physician must ultimately rely on her experience in formulating an assessment and

recommendation for her patient. Furthermore, the patient must place a great deal of trust in his physician. The physician may show the patient his PET/SPECT images and may tell the patient a numerical risk associated with various treatment options or likelihood of a particular prognosis, but the patient may very well struggle to make sense of this information. Moreover, the patient's family will likely have questions, particularly if cancer is diagnosed but the patient opts not to have surgery. The patient and/or his family members may search online for supplemental information and may become misinformed about risks of the diagnosed condition. A difficult ordeal may become more traumatic.

[0010] Thus, there remains a need for systems and methods for improved automated analysis of medical imaging studies and communication of those results, diagnoses, prognoses, treatment recommendations, and associated risks to a patient. Of particular need is an image analysis system to consistently, efficiently, and accurately detect anatomical regions, including soft tissue organs, in the entire body.

SUMMARY OF THE INVENTION

[0011] Presented herein are systems and methods that provide for automated analysis of three-dimensional (3D) medical images of a subject in order to automatically identify specific 3D volumes within the 3D images that correspond to specific anatomical regions e.g., organs and/or tissue). Notably, the image analysis approaches described herein are not limited to a single particular organ or portion of the body. Instead, they are robust and widely applicable, providing for consistent, efficient, and accurate detection of anatomical regions, including tissue and/or organs, in the entire body. In certain embodiments, the accurate identification of one or more such volumes is used to automatically determine quantitative metrics that represent uptake of radiopharmaceuticals in particular organs and/or tissue regions. These uptake metrics can be used to assess disease state in a subject, determine a prognosis for a subject, and/or determine efficacy of a treatment modality.

[0012] The capability of the approaches described herein to handle 3D images is an important advantage over certain other image analysis that only identify 2D regions in 2D images. For example, one approach relevant for cancer detection, EXINI Diagnostics AB's Bone Scan Index (BSI) software, detects regions of suspected bone cancer (see also U.S. Patent Application No. 8,855,387, issued October 7, 2014). However, the BSI analysis is carried out on two-dimensional scintigraphy images, as opposed to on three dimensional images.

[0013] Moreover, with the increase in the choices of imaging agents available to a physician to detect cancer, there remains a need for software with utility in analyzing images from any variety of imaging agents, using multiple detection modalities (SPECT/CT, PET/CT, and the like). Functional

images such as SPECT and PET provide detailed and specific information on biological processes in the body, but their potential is only realized when combined with a detailed anatomical map so that function can be localized to individual organs and structures. Although CT and MRI provide detailed anatomical information, conventional (e.g., manual) identification of organs and structures is difficult, subjective and time consuming, making certain assessment infeasible without computer support. Accordingly, by providing platform agnostic image analysis approaches that allow for accurate and robust image segmentation applicable to a variety of imaging modalities, the systems and methods described herein facilitate image analysis of particular relevance to cancer detection, diagnosis, staging, and the like.

[0014] For example, the full body segmentation approaches described herein allow for automated analysis of combinations of anatomical and functional images in order to accurately identify and grade cancerous lesions within a subject. In particular, a PET/CT composite image can be acquired for a subject following administration of a radiopharmaceutical, such as a PSMA binding agent like PyL™. The automated, machine learning-based segmentation approaches described herein are used to identify, within the CT image of the PET/CT composite, target volumes of interest (VOIs) representing target tissue regions where cancerous lesions may be found. For example, a skeletal VOI corresponding to a graphical representation of one or more bones of the subject may be identified. Once the skeletal VOI is identified in the anatomical, CT, image, it can be mapped to the PET image to identify a corresponding skeletal volume therein. The corresponding skeletal volume in the PET image is then analyzed to detect one or more localized regions of relatively high intensity, referred to as hotspots. These hotspots correspond, physically, to local regions of increased radiopharmaceutical accumulation and, accordingly, prospective cancerous lesions.

[0015] In certain embodiments, the ability to accurately and rapidly perform full body segmentation via the approaches described herein is leveraged to provide a useful and uniform scale on which to evaluate and/or measure radiopharmaceutical uptake levels in physical lesions corresponding to detected hotspots (e.g., to thereby grade expression levels of particular biomolecules, such as PSMA). In particular, in addition to detecting target VOIs corresponding to specific target tissue regions in which cancerous lesions may occur, additional target VOIs corresponding to reference tissue regions are also detected. These reference VOIs are also mapped to the PET image, to identify corresponding reference volumes therein. Measures of intensity, such as a mean, peak, maximum, etc., within these reference volumes are then computed and used as reference points against which to evaluate intensities of individual detected hotspots and convert them to index values on the scale.

[0016] For example, in certain embodiments, an aorta and a liver VOI, corresponding to a representation of a portion of an aorta and liver, respectively, are identified within the anatomical image and mapped to the functional image to identify corresponding reference volumes therein. Intensity levels of each of these reference volumes are determined (e.g., as a mean, median, peak, etc. of intensities of voxels within each reference volume), and assigned corresponding index levels on a scale. Then, for each particular individual hotspot, a hotspot intensity level is determined (e.g., similarly, as a mean, median, peak, etc. of voxel intensities within the detected hotspot). A corresponding individual hotspot index value is then determined based individual hotspot intensity level, the aorta reference intensity level, and the liver reference intensity level. This approach provides a standardized scale on which to evaluate and measure uptake associated with hotspots across different images. This allows, for example, for comparison of multiple images obtained for a single subject at different time points, as well as comparison between images of different subjects.

[0017] In certain embodiments, individual hotspot indices are used to compute an overall index for the subject and/or particular target tissue region analyzed for presence of cancerous lesions. The overall index can serve as an indicator of disease severity and/or risk, for example by reflecting a total lesion volume within a particular target tissue region, with the weighting based on the hotspot index values for the individual detected hotspots to account for radiopharmaceutical uptake within the lesions. For example, a PSMA weighted total lesion volume within a skeletal volume can be computed as a weighted sum of individual detected hotspot volumes weighted by their corresponding index values. Such an index may reflect a level and/or aggressiveness of metastasis into bone.

[0018] Volumes corresponding to other tissue regions may be similarly identified and used to determine overall index values. For example, lymph node volumes may also be used to assess severity of metastases. Volumes corresponding to tissue regions where localized disease is initially found may also be identified and analyzed to detect lesions. For example, a prostate volume may be used to assess prostate cancer severity at its initial stages. Likewise, breast volumes representing breast tissue regions and lung volumes representing lung tissue regions can be used for assessment of localized breast and lung cancers, respectively.

[0019] According, by providing for full body segmentation and automated image-based analysis of cancerous lesions throughout various relevant target tissue regions throughout a subject's body, the AI-based systems and methods described herein allow for analysis of a variety of cancers at various stages. The techniques can be used to identify and stage localized disease in relevant tissue regions such as the prostate, for example for early stage screening and monitoring, as well as to monitor regions such as bone and lymph nodes for metastasis as disease progresses. As

such, the approaches described herein provide a complete set of tools for automated image-based detection of cancer, and tracking disease evolution, progression, and response to treatment.

[0020] In one aspect, the invention is directed to a method for automatically processing a 3D image to identify 3D volumes within the 3D image that correspond to particular target tissue regions, the method comprising: (a) receiving, by a processor of a computing device, a 3D anatomical image of a subject obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within the subject; (b) automatically identifying, by the processor, using one or more machine learning modules (e.g., wherein at least one of the one or more machine learning modules is a Convolutional Neural Network (CNN) module) for each of a plurality of target tissue regions, a corresponding target volume of interest (VOI) within the 3D anatomical image; (c) determining, by the processor, a 3D segmentation map representing a plurality of 3D segmentation masks, each 3D segmentation mask representing a particular identified target VOI (e.g., automatically, digitally stitching together the plurality of 3D segmentation masks to form the 3D segmentation map); and (d) storing and/or providing for display and/or further processing, the 3D segmentation map.

[0021] In certain embodiments, the 3D anatomical image is a full body image.

[0022] In certain embodiments, step (c) comprises digitally stitching together the plurality of 3D segmentation masks to form the 3D segmentation map {e.g., by creating an initially empty image volume (e.g., initializing all voxel values to zero) and then inserting labels from each segmentation mask into the image volume [e.g., by mapping labeled (e.g., as representing a particular target tissue region as determined by a machine learning module) voxels of input images to one or machine learning modules to voxels of the image volume (e.g., so as to match voxels of the image volume to voxels of the input images that represent a same physical location, thereby labeling voxels of the image volume correctly)]}.

[0023] In certain embodiments, step (b) comprises, for at least one specific target tissue region: determining, using a first module (e.g., a localization module)(e.g., a first machine learning module), an initial VOI within the 3D anatomical image, the initial VOI corresponding to an anatomical region (e.g., a group of related tissue, such as a pelvic region, a chest region, a head and/or neck region, and the like) containing the specific target tissue region (e.g., wherein the initial VOI excludes more voxels of the 3D anatomical image than it includes; e.g., wherein the initial VOI includes less than 25% of the voxels of the 3D anatomical image; e.g., wherein a majority of voxels within the initial VOI represent physical volumes within the anatomical region); and identifying,

using a second module (e.g., a segmentation module)(e.g., a second machine learning module), the target VOI corresponding to the specific target tissue region within the initial VOI.

[0024] In certain embodiments, the second module is a CNN module that implements a CNN.

[0025] In certain embodiments, the first module is a CNN module that implements a CNN to perform a coarse segmentation to automatically identify the initial VOI corresponding to the anatomical region containing the specific target tissue region [e.g., by automatically identifying a graphical representation of a group of related tissue within the anatomical image (e.g., and, subsequently, determining the initial VOI as a rectangular region (e.g., rectangular prism or rectangular box) entirely enclosing the identified graphical representation of the group of related tissue)][e.g., by, for each of one or more particular tissue regions anticipated to be located within the anatomical region, automatically identifying a corresponding VOI in the anatomical image (e.g., via the coarse segmentation), e.g., and determining the initial VOI as a rectangular region entirely enclosing all of the identified VOIs corresponding to the particular tissue regions].

[0026] In certain embodiments, the first module receives a sub-sampled [e.g., by a factor of two or more (e.g., four) along one or more dimensions] version of the anatomical image as input and identifies the initial VOI using the sub-sampled version of the anatomical image [e.g., and wherein the second module receives a full resolution version of the anatomical image, cropped to the initial VOI (e.g., such that the first module operates on an lower resolution image that represents a larger physical volume than the second module, while the second module operates on a higher resolution image but representing a smaller physical volume)].

[0027] In certain embodiments, the first module is a first CNN module and the second module is a second CNN module, and wherein the first CNN module comprises additional filters in order to account for increased variability in image size with respect to the second CNN module.

[0028] In certain embodiments, the 3D anatomical image is a full body image and step (b) comprises: automatically determining, using one or more localization modules implementing machine learning technique(s) (e.g., wherein each localization module is a CNN module that implements a CNN), a plurality of initial VOIs within the 3D anatomical image, each initial VOI corresponding to a particular anatomical region (e.g., a group of related tissue, such as a pelvic region, a chest region, a head and/or neck region, a spine region, an upper body region, a lower body region, etc.) and in which an associated subset of the target VOIs are located; and for each initial VOI, automatically identifying, using one or more segmentation modules implementing machine learning technique(s) (e.g., wherein each segmentation module is a CNN module that implements a CNN) the associated subset of target VOIs.

[0029] In certain embodiments, the plurality of initial VOIs comprises one or more members selected from the group consisting of: a pelvic region initial VOI corresponding to a pelvic region of the subject (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of a left and/or right ilium, a sacrum, and a coccyx); a spine region initial VOI corresponding to a spine of the subject (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of a thoracic vertebra, a lumbar vertebra, and a sternum); a left upper body region initial VOI corresponding to a left side of the subject's upper body (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of one or more left rib(s), a left scapula, and a left clavicle); and a right upper body region initial VOI corresponding to a right side of the subject's upper body (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of one or more right rib(s), a right scapula, and a right clavicle).

[0030] In certain embodiments, the method comprises: (e) receiving, by the processor, a 3D functional image of the subject obtained using a functional imaging modality [e.g., single-photon emission computed tomography (SPECT); e.g., positron emission tomography (PET)][e.g., wherein the 3D functional image comprises a plurality of voxels, each representing a particular physical volume within the subject and having an intensity value that represents detected radiation emitted from a the particular physical volume, wherein at least a portion of the plurality of voxels of the 3D functional image represent physical volumes within one or more of the target tissue regions of the subject]; and (f) identifying, by the processor, within the 3D functional image, one or more 3D volume(s), each corresponding to an identified target VOI, using the 3D segmentation map (e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image).

[0031] In certain embodiments, the method comprises: (g) determining, by the processor, a cancer status [(e.g., a prostate cancer status; e.g., a metastatic cancer status (e.g., metastatic cancer, including, e.g., metastatic prostate cancer, breast cancer, lung cancer, colon cancer, skin cancer, etc.))] for the subject (e.g., using intensities of voxels within the functional image and the one or more identified 3D volumes)(e.g., based on detected lesions) [e.g., a likelihood of the subject having and/or developing prostate cancer and/or a particular stage of prostate cancer (e.g., metastatic prostate cancer)][e.g., a likelihood of the subject having and/or developing a metastatic cancer (including, e.g., metastatic prostate cancer, breast cancer, lung cancer, colon cancer, skin cancer, etc.)].

[0032] In certain embodiments, the method comprises performing steps (a) – (g) repeatedly for a plurality of anatomical and corresponding functional images collected at different time points to determine, at each time point, a cancer status of the subject, thereby tracking cancer status over time (e.g., to evaluate disease progression and/or treatment efficacy).

[0033] In certain embodiments, the method comprises (e.g., prior to step (g)): (h) automatically adjusting, by the processor, intensities of voxels of the 3D functional image to correct for background uptake (e.g., of the radiopharmaceutical) in one or more background tissue regions (e.g., correcting for uptake of the radiopharmaceutical that occurs in the one or more background tissue regions under normal circumstances, and which is not necessarily indicative of presence of cancerous lesions).

[0034] In certain embodiments, the target tissue regions comprise the one or more background tissue regions and step (h) comprises: using the 3D segmentation map to identify, within the 3D functional image, one or more 3D background tissue volume(s), each corresponding a particular background tissue region (e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image); and adjusting intensities of voxels of the 3D functional image based on intensities of voxels within the 3D background tissue volumes (e.g., by using intensities of voxels within the 3D background tissue volumes to estimate a contribution to intensities outside of the 3D background tissue volumes).

[0035] In certain embodiments, the one or more background tissue regions comprise one or more members selected from the group consisting of: a bladder (e.g., a urinary bladder), a kidney, a duodenum, a small intestine, a spleen, a liver, a pancreas, a stomach, an adrenal gland, a rectum, and testes.

[0036] In certain embodiments, the method comprises: (i) automatically detecting, by the processor, one or more hotspots within the 3D functional image determined to represent lesions based on intensities of voxels within the 3D functional image [e.g., based on a comparison of intensities within the 3D functional image with a threshold value (e.g., wherein the 3D functional image is a 3D PET image and the threshold is a particular standard uptake value (SUV) level)] (e.g., and also based on the one or more 3D volumes identified within the 3D functional image).

[0037] In certain embodiments, step (i) comprises using one or more thresholds [e.g., comparing intensities of voxels in the 3D functional image with the one or more thresholds (e.g., wherein the one or more thresholds comprises a plurality of region-specific thresholds, each used for a specific sub-set of the one or more 3D volumes identified within the 3D functional image)].

[0038] In certain embodiments, step (i) comprises applying one or more filters to the 3D functional image [e.g., as in a blob detection technique; e.g., wherein the one or more filters

comprises one or more Gaussian filters (e.g., a high pass Gaussian filter and a low pass Gaussian filter, as in a Difference of Gaussians approach); e.g., wherein the one or more filters comprises a Laplacian filter (e.g., as in a Laplacian of a Gaussian technique); e.g., wherein different filter kernels are used for different sub-sets of the one or more 3D volumes identified within the 3D functional image].

[0039] In certain embodiments, step (i) comprises using a combination of two or more techniques for identification of hotspots in different sub-sets of the one or more 3D volumes identified within the 3D functional image [e.g., using a first filtering approach (e.g., as in paragraph [0038]) for a first sub-set of the one or more 3D volumes and using a second filtering approach for a second sub-set; e.g., using thresholding approach (e.g., as paragraph [0037]) for a first sub-set of the one or more 3D volumes and using a filtering approach (e.g., as paragraph [0038]) for a second sub-set of the one or more 3D volumes].

[0040] In certain embodiments, step (i) comprises detecting an initial set of hotspots and, for at least a portion of the hotspots of the initial set, classifying each hotspot of at least a portion of the detected hotspots as either a cancerous lesion or not a cancerous lesion (e.g., as noise)[e.g., using a machine learning module; e.g., based on a shape and/or location of the hotspot (e.g., in combination with anatomical knowledge; e.g., wherein the location includes an identification of a particular target tissue region corresponding to the 3D volume in which the hotspot is located and/or a relative position of the hotspot within the particular target tissue region); e.g., and removing hotspots classified as not a cancerous lesion from the initial set, thereby obtaining a final set of hotspots determined to represent lesions].

[0041] In certain embodiments, the target tissue regions comprise one or more background tissue regions and the method comprises: using the 3D segmentation map to identify, within the 3D functional image, one or more 3D background tissue volume(s), each corresponding a particular background tissue region (e.g., a background tissue region in which significant radiopharmaceutical uptake occurs under normal circumstances and is not necessarily indicative of presence of cancerous lesions)(e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image); and excluding voxels of the 3D within the 3D background tissue from the voxels used to automatically detect the one or more hotspots at step (i).

[0042] In certain embodiments, the method comprises using the one or more detected hotspots (e.g., intensities of the one or more detected hotspots) to determine a cancer status for the subject.

[0043] In certain embodiments, the target tissue regions comprise one or more reference tissue regions and the method comprises: using the 3D segmentation map to identify, by the

processor, within the 3D functional image, one or more 3D reference volume(s), each corresponding to a particular reference tissue region; determining, by the processor, one or more reference intensity values, each associated with a particular 3D reference volume of the one or more 3D reference volume(s) and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity within the particular 3D reference volume; determining, by the processor, one or more individual hotspot intensity values, each associated with a particular hotspot of at least a portion of the detected one or more hotspots and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity of the particular hotspot; and determining, by the processor, one or more individual hotspot index values using the one or more individual hotspot intensity values and the one or more reference intensity values [e.g., wherein each individual hotspot index value is associated with a particular hotspot of the at least a portion of the detected one or more hotspots and determined using (e.g., based on a comparison between) (i) the individual hotspot intensity value associated with the particular hotspot and (ii) the one or more reference intensity values].

[0044] In certain embodiments, the determining the one or more individual hotspot index values comprises mapping each of the one or more reference intensity values to a corresponding reference index value on a scale and, for each individual hotspot intensity value, using the reference intensity values and corresponding reference index values to interpolate a corresponding individual hotspot index value.

[0045] In certain embodiments, the reference tissue regions comprise one or more members selected from the group consisting of: a liver, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), and a parotid gland.

[0046] In certain embodiments, a first reference intensity value (i) is a blood reference intensity value associated with a 3D reference volume corresponding to an aorta portion, and (ii) maps to a first reference index value; a second reference intensity value (i) is a liver reference intensity value associated with a 3D reference volume corresponding to a liver, and (ii) maps to a second reference index value; and the second reference intensity value is greater than the first reference intensity value and the second reference index value is greater than the first reference index value.

[0047] In certain embodiments, the reference intensity values comprises a maximum reference intensity value that maps to a maximum reference index value, and hotspots having associated hotspot intensity values greater than the maximum reference intensity value are assigned hotspot index values equal to the maximum reference index value.

[0048] In certain embodiments, the method comprises determining, by the processor, an overall index value indicative of a cancer status of the subject using at least a portion of the one or more hotspot index values.

[0049] In certain embodiments, the overall index value is determined as a weighted sum of at least a portion [e.g., located within a 3D volume of the functional image that corresponds to a skeletal region of the subject; e.g., located within a 3D volume of the functional image that corresponds to a lymph region of the subject] of the individual hotspot index values [e.g., wherein each hotspot index value in the sum is weighted by a measure of size (e.g., 3D volume; e.g., average diameter) of the associated hotspot].

[0050] In certain embodiments, the overall index value is associated with a particular target tissue region corresponding to a particular target VOI identified within the anatomical image; and the overall index value is determined using hotspot index values of a subset of hotspots located within a particular 3D volume in the 3D functional image that corresponds to the particular identified target VOI [e.g., and wherein the overall index value is computed as a weighted sum of the hotspot index values of the subset of hotspots (e.g., wherein the weighted sum is normalized by an estimated volume of the particular target tissue region (e.g., computed as a volume of the particular 3D volume in the function image and/or a volume of the particular target VOI))].

[0051] In certain embodiments, the particular target tissue region is selected from the group consisting of: a skeletal region comprising one or more bones of the subject, a lymph region, and a prostate region.

[0052] In certain embodiments, the 3D anatomical image is an x-ray computed tomography (CT) image, and the 3D functional image is a 3D single photon emission computed tomography (SPECT) image.

[0053] In certain embodiments, the 3D anatomical image is an x-ray computed tomography (CT) image, and the 3D functional image is a 3D positron emission tomography (PET) image.

[0054] In certain embodiments, the 3D PET image of the subject is obtained following administration to the subject of a radiopharmaceutical comprising a prostate-specific membrane antigen (PSMA) binding agent.

[0055] In certain embodiments, the radiopharmaceutical comprises [18F]DCFPyL.

[0056] In certain embodiments, step (b) comprises cropping, by the processor, the 3D anatomical image to remove voxels representing air [e.g., to create a cropped anatomical image, and using the cropped anatomical image to identify the one or more target VOIs (e.g., using the cropped anatomical image is used as input to one or more machine learning modules, as opposed to the original size 3D anatomical image)].

[0057] In certain embodiments, the target tissue regions comprise one or more members selected from the group consisting of: a left hip bone, a right hip bone, a sacrum and coccyx region, a left clavicle, a right clavicle, a left rib, a right rib, a left scapula, a right scapula, a sternum, a lumbar vertebra, a thoracic vertebra, a skull, a cervical vertebra, a left femur, a right femur, a left humerus, a right humerus, a prostate, a urinary bladder, a rectum, a left gluteus maximus, a right gluteus maximus, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), a left kidney, a right kidney, a liver, a left lung, a right lung, a spleen, a ventricle, a left adrenal gland, a right adrenal gland, a gallbladder, a brain, a pancreas, a heart, a mandible, a left bronchi, a right bronchi, a trachea, a left common iliac artery, a right common iliac artery, and a parotid gland.

[0058] In certain embodiments, one or more of the target tissue regions comprise one or more specific bones selected from the group consisting of: a left clavicle, a right clavicle, a left femur, a right femur, a left fibula, a right fibula, a left hip bone, a right hip bone, a left humerus, a right humerus, a mandible, a left patella, a right patella, left radius, a right radius, a left tibia, a right tibia, a left ulna, a right ulna, a left rib (e.g., a first left rib, a second left rib, a third left rib, a fourth left rib, a fifth left rib, a sixth left rib, a seventh left rib, an eighth left rib, a ninth left rib, a tenth left rib, an eleventh left rib, a twelfth left rib), a right rib (e.g., a first right rib, a second right rib, a third right rib, a fourth right rib, a fifth right rib, a sixth right rib, a seventh right rib, an eighth right rib, a ninth right rib, a tenth right rib, an eleventh right rib, a twelfth right rib), a sacrum and coccyx (e.g., a combined sacrum and coccyx region; e.g., a sacrum and a coccyx individually, so as to distinguish between the two), a left scapula, a right scapula, a skull, sternum, a vertebrae region [e.g., a cervical vertebrae region, comprising one or more (e.g., up to all) cervical vertebrae; e.g., a lumbar vertebrae region, comprising one or more (e.g., up to all) lumbar vertebrae; e.g., a thoracic vertebrae region, comprising one or more (e.g., up to all) thoracic vertebrae], and an individual vertebra [e.g., e.g., an individual cervical vertebra (e.g., a first cervical vertebra, a second cervical vertebra, a third cervical vertebra, a fourth cervical vertebra, a fifth cervical vertebra, a sixth cervical vertebra, a seventh cervical vertebra); e.g., an individual lumbar vertebra (e.g., a first lumbar vertebra, a second lumbar vertebra, a third lumbar vertebra, a fourth lumbar vertebra, a fifth lumbar vertebra, a sixth lumbar vertebra); e.g., an individual thoracic vertebra (e.g., a first thoracic vertebra, a second thoracic vertebra, a third thoracic vertebra, a fourth, thoracic vertebra, a fifth thoracic vertebra, a sixth thoracic vertebra, a seventh thoracic vertebra, an eighth thoracic vertebra, a ninth thoracic vertebra, a tenth thoracic vertebra, an eleventh thoracic vertebra, an twelfth thoracic vertebra)].

[0059] In certain embodiments, one or more of the target tissue regions comprise one or more specific bones selected from the group consisting of: a left clavicle, a right clavicle, a left hip bone, a right hip bone, a left rib (e.g., a first left rib, a second left rib, a third left rib, a fourth left rib,

a fifth left rib, a sixth left rib, a seventh left rib, an eighth left rib, a ninth left rib, a tenth left rib, an eleventh left rib, a twelfth left rib), a right rib (e.g., a first right rib, a second right rib, a third right rib, a fourth right rib, a fifth right rib, a sixth right rib, a seventh right rib, an eighth right rib, a ninth right rib, a tenth right rib, an eleventh right rib, a twelfth right rib), a sacrum and coccyx (e.g., a combined sacrum and coccyx region; e.g., a sacrum and a coccyx individually, so as to distinguish between the two), a left scapula, a right scapula, a sternum, a vertebrae region [e.g., a cervical vertebrae region, comprising one or more (e.g., up to all) cervical vertebrae; e.g., a lumbar vertebrae region, comprising one or more (e.g., up to all) lumbar vertebrae; e.g., a thoracic vertebrae region, comprising one or more (e.g., up to all) thoracic vertebrae], and an individual vertebra [e.g., an individual lumbar vertebra (e.g., a first lumbar vertebra, a second lumbar vertebra, a third lumbar vertebra, a fourth lumbar vertebra, a fifth lumbar vertebra); e.g., an individual thoracic vertebra (e.g., a first thoracic vertebra, a second thoracic vertebra, a third thoracic vertebra, a fourth, thoracic vertebra, a fifth thoracic vertebra, a sixth thoracic vertebra, a seventh thoracic vertebra, an eighth thoracic vertebra, a ninth thoracic vertebra, a tenth thoracic vertebra, an eleventh thoracic vertebra, an twelfth thoracic vertebra)].

[0060] In certain embodiments, one or more of the target tissue regions comprise one or more soft-tissue regions (e.g., organs) selected from the group consisting of: a left adrenal gland, a right adrenal gland, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), a brain, a left bronchi, a right bronchi, a gallbladder, a left gluteus maximus, a right gluteus maximus, a heart, a left common iliac artery, a right common iliac artery, a left kidney, a right kidney, a liver, a left lung, a right lung, a pancreas, a prostate, a rectum, a spleen , a trachea, a urinary bladder, a ventricle, and a parotid gland.

[0061] In certain embodiments, the one or more target tissue regions comprise one or more soft-tissue regions (e.g., organs) selected from the group consisting of: a gallbladder, a left kidney, a right kidney, a liver, a left lung, a right lung, a prostate, and a urinary bladder.

[0062] In certain embodiments, the target tissue regions comprise one or more bone regions, each corresponding to a specific bone [e.g., where the target tissue regions comprise at least a portion of (e.g., up to all) the bones listed in paragraph [0058] or [0059]] and one or more soft tissue regions, each corresponding to a specific soft-tissue region [e.g., where the target tissue regions comprise at least a portion of (e.g., up to all) the soft-tissue regions listed in paragraph [0060] or [0061]].

[0063] In another aspect, the invention is directed to a method for automatically processing 3D images to automatically identify cancerous lesions within a subject, the method comprising: (a) receiving, by a processor of a computing device, a 3D anatomical image of a subject obtained using

an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within the subject; (b) automatically identifying, by the processor, using one or more machine learning modules (e.g., wherein at least one of the one or more machine learning modules is a Convolutional Neural Network (CNN) module), for each of a plurality of target tissue regions, a corresponding target volume of interest (VOI) within the 3D anatomical image; (c) determining, by the processor, a 3D segmentation map representing a plurality of 3D segmentation masks, each 3D segmentation mask representing a particular identified target VOI (e.g., automatically, digitally stitching together the plurality of 3D segmentation masks to form the 3D segmentation map); (d) receiving, by the processor, a 3D functional image of the subject obtained using a functional imaging modality [e.g., single-photon emission computed tomography (SPECT); e.g., positron emission tomography (PET)][e.g., wherein the 3D functional image comprises a plurality of voxels, each representing a particular physical volume within the subject and having an intensity value that represents detected radiation emitted from a the particular physical volume, wherein at least a portion of the plurality of voxels of the 3D functional image represent physical volumes within one or more of the target tissue regions of the subject]; (e) identifying, within the 3D functional image, one or more 3D volume(s), each corresponding to an identified target VOI, using the 3D segmentation map (e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image); and (f) automatically detecting, by the processor, within at least a portion of the one or more 3D volumes identified within the 3D functional image, one or more hotspots determined to represent lesions based on intensities of voxels within the 3D functional image [e.g., based on a comparison of intensities within the 3D functional image with a threshold value (e.g., wherein the 3D functional image is a 3D PET image and the threshold is a particular standard uptake value (SUV) level)].

[0064] In certain embodiments, the method comprises using, by the processor, the one or more detected hotspots (e.g., intensities of the one or more detected hotspots) to determine a cancer status for the subject.

[0065] In certain embodiments, the target tissue regions comprise one or more reference tissue regions and wherein the method comprises: using the 3D segmentation map to identify, by the processor, within the 3D functional image, one or more 3D reference volume(s), each corresponding to a particular reference tissue region; determining one or more reference intensity values, each associated with a particular 3D reference volume of the one or more 3D reference volume(s) and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity within the particular 3D reference volume; determining, by the processor, one or more

individual hotspot intensity values, each associated with a particular hotspot of at least a portion of the detected one or more hotspots and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity of the particular hotspot; and determining, by the processor, one or more individual hotspot index values using the one or more individual hotspot intensity values and the one or more reference intensity values [e.g., wherein each individual hotspot index value is associated with a particular hotspot of the at least a portion of the detected one or more hotspots and determined using (e.g., based on a comparison between) (i) the individual hotspot intensity value associated with the particular hotspot and (ii) the one or more reference intensity values].

[0066] In certain embodiments, the reference tissue regions comprise one or more members selected from the group consisting of: a liver, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), and a parotid gland.

[0067] In certain embodiments, the method comprises determining, by the processor, an overall index value indicative of a cancer status of the subject using at least a portion of the one or more hotspot index values.

[0068] In certain embodiments, the overall index value is determined as a weighted sum of at least a portion [e.g., located within a 3D volume of the functional image that corresponds to a skeletal region of the subject; e.g., located within a 3D volume of the functional image that corresponds to a lymph region of the subject] of the individual hotspot index values [e.g., wherein each hotspot index value in the sum is weighted by a measure of size (e.g., 3D volume; e.g., average diameter) of the associated hotspot].

[0069] In certain embodiments, the 3D anatomical image is an x-ray computed tomography (CT) image, and the 3D functional image is a 3D positron emission tomography (PET) image.

[0070] In certain embodiments, the 3D PET image of the subject is obtained following administration to the subject of a radiopharmaceutical comprising a prostate-specific membrane antigen (PSMA) binding agent.

[0071] In certain embodiments, the radiopharmaceutical comprises [18F]DCFPyL.

[0072] In another aspect, the invention is directed to a method for automatically processing 3D images to identify, and measure uptake of radiopharmaceutical in, cancerous lesions (e.g., metastases) within a subject having or at risk for a cancer (e.g., prostate cancer, breast cancer, lung cancer; e.g., a metastatic cancer, such as metastatic prostate cancer, metastatic breast cancer, metastatic lung cancer), the method comprising: (a) receiving, by a processor of a computing device, a 3D anatomical image of a subject obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone)

within the subject; (b) automatically identifying, by the processor, using one or more machine learning modules (e.g., wherein at least one of the machine learning modules is a Convolutional Neural Network (CNN) module), within the 3D anatomical image: a first skeletal volume comprising a graphical representation of one or more bones of the subject; a first aorta volume comprising a graphical representation of at least a portion of an aorta of the subject; and a first liver volume comprising a graphical representation of a liver of the subject; (c) determining, by the processor, a 3D segmentation map representing a plurality of 3D segmentation masks, including a skeletal mask representing the identified first skeletal volume, an aorta mask representing the identified first aorta volume, and a liver mask representing the identified first liver volume; (d) receiving, by the processor, a 3D functional image of the subject obtained using a functional imaging modality [e.g., single-photon emission computed tomography (SPECT); e.g., positron emission tomography (PET)][e.g., wherein the 3D functional image comprises a plurality of voxels, each representing a particular physical volume within the subject and having an intensity value that represents detected radiation emitted from a the particular physical volume, wherein at least a portion of the plurality of voxels of the 3D functional image represent physical volumes within one or more bones, an aorta portion, and/or a liver of the subject]; (e) automatically identifying, within the 3D functional image, using the 3D segmentation map (e.g., by mapping the 3D segmentation masks of the 3D segmentation map to the 3D functional image): a second skeletal volume corresponding to the first identified skeletal volume, within the 3D anatomical image; a second aorta volume corresponding to the first aorta volume, identified within the 3D anatomical image; and a second liver volume corresponding to the first liver volume, identified within the 3D anatomical image; (f) automatically detecting, by the processor, within the second skeletal volume, one or more hotspots determined to represent lesions based on intensities of voxels within the second skeletal volume (e.g., the one or more hotspots corresponding to localized regions of relatively high intensity, e.g., identified based on a comparison of intensities of voxels located within the second skeletal volume with a threshold value); and (g) determining, by the processor, for each of the one or more detected hotspots, an individual hotspot index (e.g., indicative of a measure of radiopharmaceutical uptake in the lesions represented by the detected hotspot) value by: determining an aorta reference intensity level based on a measure (e.g., a mean, a maximum, a median, etc.) of intensity of voxels within the second aorta volume; determining a liver reference intensity level based on a measure (e.g., a mean, a maximum, a median, etc.) of intensity of voxels within the second liver volume; and for each individual detected hotspot: determining a corresponding individual hotspot intensity level based on a measure (e.g., a mean, a maximum, a median, etc.) of intensity of voxels of the detected hotspot;

and determining a corresponding individual hotspot index level from the individual hotspot intensity level, the aorta reference intensity level, and the liver reference intensity level.

[0073] In certain embodiments, the method comprises determining, by the processor, an overall index value indicative of a cancer status of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots.

[0074] In certain embodiments, the subject has or is at risk for prostate cancer.

[0075] In certain embodiments, step (b) comprises automatically identifying, within the 3D anatomical image, a first prostate volume comprising a graphical representation of a prostate of the subject, the 3D segmentation map determined at step (c) further includes a prostate mask representing the identified first prostate volume, step (e) comprises automatically identifying, within the 3D functional image, a second prostate volume corresponding to the first identified prostate volume, within the 3D anatomical image, step (f) comprises automatically detecting one or more hotspots in the second prostate volume, and the method further comprises: determining, by the processor, (i) an overall bone index value indicative of a lesion content (e.g., and severity) in the one or more bones of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second skeletal volume and (ii) an overall prostate index value indicative of a lesion content (e.g., and severity) in the prostate of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second prostate volume.

[0076] In certain embodiments, the subject has or is at risk for metastatic cancer (e.g., metastatic prostate cancer, metastatic breast cancer, metastatic lung cancer, and other metastatic bone cancers).

[0077] In another aspect, the invention is directed to a method for automatically processing a set of full body 3D anatomical images of varying sizes to automatically identify, within each 3D anatomical image, a plurality of 3D volumes that correspond to particular target tissue regions, the method comprising: (a) receiving, by a processor of a computing device, the set of 3D anatomical images of one or more subject(s) obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within each of the one or more subject(s), wherein the set of 3D anatomical images has a mean x-dimension, a mean y-dimension, and a mean z-dimension, at least one of which (the mean x-, mean y-, or mean z-dimension) having a standard deviation at least 3% of the corresponding mean [e.g., at least one of which (the mean x-, mean y-, or mean z-dimension) having a relative standard deviation of at least 3 %, at least 5%, at least 10%, at least 15%, or at least 20%]; and (b) automatically

determining, by the processor, using a localization module that implements a CNN, within each image of the set of 3D anatomical images, at least one initial VOI corresponding to a particular anatomical region (e.g., a group of related tissue, such as a pelvic region and a spin region) that comprises one or more particular associated target tissue regions (e.g., one or more specific bones, e.g., one or more specific organs), thereby identifying at least one initial VOI for the corresponding 3D anatomical image; and (c) for each initial VOI, automatically identifying, by the processor, using one or more segmentation modules each implementing a CNN, for each of the one or more particular target tissue regions associated with the particular anatomical region to which the initial VOI corresponds, a corresponding target VOI (e.g., a graphical representation of the particular target tissue region).

[0078] In another aspect, the invention is directed to a system for automatically processing a 3D image to identify 3D volumes within the 3D image that correspond to particular target tissue regions, the system comprising: a processor of a computing device; and a memory having instructions stored thereon, wherein the instructions, when executed by the processor, cause the processor to: (a) receive a 3D anatomical image of a subject obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultrasound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within the subject; (b) automatically identify, using one or more machine learning modules (e.g., wherein at least one of the one or more machine learning modules is a Convolutional Neural Network (CNN) module), for each of a plurality of target tissue regions, a corresponding target volume of interest (VOI) within the 3D anatomical image; (c) determine 3D segmentation map representing a plurality of 3D segmentation masks, each 3D segmentation mask representing a particular identified target VOI (e.g., automatically, digitally stitching together the plurality of 3D segmentation masks to form the 3D segmentation map); and (d) store and/or provide the 3D segmentation map for display and/or further processing.

[0079] In certain embodiments, the 3D anatomical image is a full body image.

[0080] In certain embodiments, at step (c) the instructions cause the processor to digitally stitch together the plurality of 3D segmentation masks to form the 3D segmentation map {e.g., by creating an initially empty image volume (e.g., initializing all voxel values to zero) and then inserting labels from each segmentation mask into the image volume [e.g., by mapping labeled (e.g., as representing a particular target tissue region as determined by a machine learning module) voxels of input images to one or machine learning modules to voxels of the image volume (e.g., so as to match voxels of the image volume to voxels of the input images that represent a same physical location, thereby labeling voxels of the image volume correctly)]}.

[0081] In certain embodiments, at step (b), the instructions cause the processor to, for at least one specific target tissue region: determine, using a first module (e.g., a localization module)(e.g., a first machine learning module), an initial VOI within the 3D anatomical image, the initial VOI corresponding to an anatomical region (e.g., a group of related tissue, such as a pelvic region, a chest region, a head and/or neck region, and the like) containing the specific target tissue region (e.g., wherein the initial VOI excludes more voxels of the 3D anatomical image than it includes; e.g., wherein the initial VOI includes less than 25% of the voxels of the 3D anatomical image; e.g., wherein a majority of voxels within the initial VOI represent physical volumes within the anatomical region); and identify, using a second module (e.g., a segmentation module)(e.g., a second machine learning module), the target VOI corresponding to the specific target tissue region within the initial VOI.

[0082] In certain embodiments, the second module is a CNN module that implements a CNN.

[0083] In certain embodiments, the first module is a CNN module that implements a CNN to perform a coarse segmentation to automatically identify the initial VOI corresponding to the anatomical region containing the specific target tissue region [e.g., by automatically identifying a graphical representation of a group of related tissue within the anatomical image (e.g., and, subsequently, determining the initial VOI as a rectangular region (e.g., rectangular prism or rectangular box) entirely enclosing the identified graphical representation of the group of related tissue)][e.g., by, for each of one or more particular tissue regions anticipated to be located within the anatomical region, automatically identifying a corresponding VOI in the anatomical image (e.g., via the coarse segmentation), e.g., and determining the initial VOI as a rectangular region entirely enclosing all of the identified VOIs corresponding to the particular tissue regions].

[0084] In certain embodiments, the first module receives a sub-sampled [e.g., by a factor of two or more (e.g., four) along one or more dimensions] version of the anatomical image as input and identifies the initial VOI using the sub-sampled version of the anatomical image [e.g., and wherein the second module receives a full resolution version of the anatomical image, cropped to the initial VOI (e.g., such that the first module operates on an lower resolution image that represents a larger physical volume than the second module, while the second module operates on a higher resolution image but representing a smaller physical volume)].

[0085] In certain embodiments, the first module is a first CNN module and the second module is a second CNN module, and wherein the first CNN module comprises additional filters in order to account for increased variability in image size with respect to the second CNN module.

[0086] In certain embodiments, the 3D anatomical image is a full body image and wherein, at step (b), the instructions cause the processor to: automatically determine, using one or more localization modules implementing machine learning technique(s) (e.g., wherein each localization module is a CNN module that implements a CNN), a plurality of initial VOIs within the 3D anatomical image, each initial VOI corresponding to a particular anatomical region (e.g., a group of related tissue, such as a pelvic region, a chest region, a head and/or neck region, a spine region, an upper body region, a lower body region, etc.) and in which an associated subset of the target VOIs are located; and for each initial VOI, automatically identify, using one or more segmentation modules implementing machine learning technique(s) (e.g., wherein each segmentation module is a CNN module that implements a CNN) the associated subset of target VOIs.

[0087] In certain embodiments, the plurality of initial VOIs comprises one or more members selected from the group consisting of: a pelvic region initial VOI corresponding to a pelvic region of the subject (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of a left and/or right ilium, a sacrum, and a coccyx); a spine region initial VOI corresponding to a spine of the subject (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of a thoracic vertebra, a lumbar vertebra, and a sternum); a left upper body region initial VOI corresponding to a left side of the subject's upper body (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of one or more left rib(s), a left scapula, and a left clavicle); and a right upper body region initial VOI corresponding to a right side of the subject's upper body (e.g., wherein the subset of target VOIs located in the pelvic region initial VOI comprise one or more target VOIs corresponding to target tissue regions selected from the group consisting of one or more right rib(s), a right scapula, and a right clavicle).

[0088] In certain embodiments, the instructions cause the processor to: (e) receive a 3D functional image of the subject obtained using a functional imaging modality [e.g., single-photon emission computed tomography (SPECT); e.g., positron emission tomography (PET)][e.g., wherein the 3D functional image comprises a plurality of voxels, each representing a particular physical volume within the subject and having an intensity value that represents detected radiation emitted from a the particular physical volume, wherein at least a portion of the plurality of voxels of the 3D functional image represent physical volumes within one or more of the target tissue regions of the subject]; and (f) identify, within the 3D functional image, one or more 3D volume(s), each

corresponding to an identified target VOI, using the 3D segmentation map (e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image).

[0089] In certain embodiments, the instructions cause the processor to: (g) determine a cancer status [(e.g., a prostate cancer status; e.g., a metastatic cancer status (e.g., metastatic cancer, including, e.g., metastatic prostate cancer, breast cancer, lung cancer, colon cancer, skin cancer, etc.)) for the subject (e.g., using intensities of voxels within the functional image and the one or more identified 3D volumes)(e.g., based on detected lesions) [e.g., a likelihood of the subject having and/or developing prostate cancer and/or a particular stage of prostate cancer (e.g., metastatic prostate cancer)][e.g., a likelihood of the subject having and/or developing a metastatic cancer (including, e.g., metastatic prostate cancer, breast cancer, lung cancer, colon cancer, skin cancer, etc.)].

[0090] In certain embodiments, the instructions cause the processor to perform steps (a) – (g) repeatedly for a plurality of anatomical and corresponding functional images collected at different time points to determine, at each time point, a cancer status of the subject, thereby tracking cancer status over time (e.g., to evaluate disease progression and/or treatment efficacy).

[0091] In certain embodiments, the instructions cause the processor to (e.g., prior to step (g)): (h) automatically adjust intensities of voxels of the 3D functional image to correct for background uptake (e.g., of the radiopharmaceutical) in one or more background tissue regions (e.g., correcting for uptake of the radiopharmaceutical that occurs in the one or more background tissue regions under normal circumstances, and which is not necessarily indicative of presence of cancerous lesions).

[0092] In certain embodiments, the target tissue regions comprise the one or more background tissue regions and wherein, at step (h), the instructions cause the processor to: use the 3D segmentation map to identify, within the 3D functional image, one or more 3D background tissue volume(s), each corresponding a particular background tissue region (e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image); and adjust intensities of voxels of the 3D functional image based on intensities of voxels within the 3D background tissue volumes (e.g., by using intensities of voxels within the 3D background tissue volumes to estimate a contribution to intensities outside of the 3D background tissue volumes).

[0093] In certain embodiments, the one or more background tissue regions comprise one or more members selected from the group consisting of: a bladder (e.g., a urinary bladder), a kidney, a duodenum, a small intestine, a spleen, a liver, a pancreas, a stomach, an adrenal gland, a rectum, and testes.

[0094] In certain embodiments, the instructions cause the processor to: (i) automatically detect one or more hotspots within the 3D functional image determined to represent lesions based on intensities of voxels within the 3D functional image [e.g., based on a comparison of intensities within the 3D functional image with a threshold value (e.g., wherein the 3D functional image is a 3D PET image and the threshold is a particular standard uptake value (SUV) level)] (e.g., and also based on the one or more 3D volumes identified within the 3D functional image).

[0095] In certain embodiments, at step (i), the instructions cause the processor to use one or more thresholds [e.g., comparing intensities of voxels in the 3D functional image with the one or more thresholds (e.g., wherein the one or more thresholds comprises a plurality of region-specific thresholds, each used for a specific sub-set of the one or more 3D volumes identified within the 3D functional image)].

[0096] In certain embodiments, at step (i), the instructions cause the processor to apply one or more filters to the 3D functional image [e.g., as in a blob detection technique; e.g., wherein the one or more filters comprises one or more Gaussian filters (e.g., a high pass Gaussian filter and a low pass Gaussian filter, as in a Difference of Gaussians approach); e.g., wherein the one or more filters comprises a Laplacian filter (e.g., as in a Laplacian of a Gaussian technique); e.g., wherein different filter kernels are used for different sub-sets of the one or more 3D volumes identified within the 3D functional image].

[0097] In certain embodiments, at step (i), the instructions cause the processor to use a combination of two or more techniques for identification of hotspots in different sub-sets of the one or more 3D volumes identified within the 3D functional image [e.g., using a first filtering approach (e.g., as in paragraph [0096]) for a first sub-set of the one or more 3D volumes and using a second filtering approach for a second sub-set; e.g., using thresholding approach (e.g., as in paragraph [0095]) for a first sub-set of the one or more 3D volumes and using a filtering approach (e.g., as in paragraph [0096]) for a second sub-set of the one or more 3D volumes].

[0098] In certain embodiments, at step (i), the instructions cause the processor to detect an initial set of hotspots and, for at least a portion of the hotspots of the initial set, classify each hotspot of at least a portion of the detected hotspots as either a cancerous lesion or not a cancerous lesion (e.g., as noise)[e.g., using a machine learning module; e.g., based on a shape and/or location of the hotspot (e.g., in combination with anatomical knowledge; e.g., wherein the location includes an identification of a particular target tissue region corresponding to the 3D volume in which the hotspot is located and/or a relative position of the hotspot within the particular target tissue region); e.g., and removing hotspots classified as not a cancerous lesion from the initial set, thereby obtaining a final set of hotspots determined to represent lesions].

[0099] In certain embodiments, the target tissue regions comprise one or more background tissue regions and the instructions cause the processor to: use the 3D segmentation map to identify, within the 3D functional image, one or more 3D background tissue volume(s), each corresponding a particular background tissue region (e.g., a background tissue region in which significant radiopharmaceutical uptake occurs under normal circumstances and is not necessarily indicative of presence of cancerous lesions)(e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image); and exclude voxels of the 3D within the 3D background tissue from the voxels used to automatically detect the one or more hotspots at step (i).

[0100] In certain embodiments, the instructions cause the processor to use the one or more detected hotspots (e.g., intensities of the one or more detected hotspots) to determine a cancer status for the subject.

[0101] In certain embodiments, the target tissue regions comprise one or more reference tissue regions and the instructions cause the processor to: use the 3D segmentation map to identify, within the 3D functional image, one or more 3D reference volume(s), each corresponding to a particular reference tissue region; determine one or more reference intensity values, each associated with a particular 3D reference volume of the one or more 3D reference volume(s) and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity within the particular 3D reference volume; determine one or more individual hotspot intensity values, each associated with a particular hotspot of at least a portion of the detected one or more hotspots and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity of the particular hotspot; and determine one or more individual hotspot index values using the one or more individual hotspot intensity values and the one or more reference intensity values [e.g., wherein each individual hotspot index value is associated with a particular hotspot of the at least a portion of the detected one or more hotspots and determined using (e.g., based on a comparison between) (i) the individual hotspot intensity value associated with the particular hotspot and (ii) the one or more reference intensity values].

[0102] In certain embodiments, the instructions cause the processor to determine the one or more individual hotspot index values by mapping each of the one or more reference intensity values to a corresponding reference index value on a scale and, for each individual hotspot intensity value, using the reference intensity values and corresponding reference index values to interpolate a corresponding individual hotspot index value.

[0103] In certain embodiments, the reference tissue regions comprise one or more members selected from the group consisting of: a liver, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), and a parotid gland.

[0104] In certain embodiments, a first reference intensity value (i) is a blood reference intensity value associated with a 3D reference volume corresponding to an aorta portion, and (ii) maps to a first reference index value; a second reference intensity value (i) is a liver reference intensity value associated with a 3D reference volume corresponding to a liver, and (ii) maps to a second reference index value; and the second reference intensity value is greater than the first reference intensity value and the second reference index value is greater than the first reference index value.

[0105] In certain embodiments, the reference intensity values comprises a maximum reference intensity value that maps to a maximum reference index value, and wherein hotspots having associated hotspot intensity values greater than the maximum reference intensity value are assigned hotspot index values equal to the maximum reference index value.

[0106] In certain embodiments, the instructions cause the processor to determine an overall index value indicative of a cancer status of the subject using at least a portion of the one or more hotspot index values.

[0107] In certain embodiments, the overall index value is determined as a weighted sum of at least a portion [e.g., located within a 3D volume of the functional image that corresponds to a skeletal region of the subject; e.g., located within a 3D volume of the functional image that corresponds to a lymph region of the subject] of the individual hotspot index values [e.g., wherein each hotspot index value in the sum is weighted by a measure of size (e.g., 3D volume; e.g., average diameter) of the associated hotspot].

[0108] In certain embodiments, the overall index value is associated with a particular target tissue region corresponding to a particular target VOI identified within the anatomical image; and the overall index value is determined using hotspot index values of a subset of hotspots located within a particular 3D volume in the 3D functional image that corresponds to the particular identified target VOI [e.g., and wherein the overall index value is computed as a weighted sum of the hotspot index values of the subset of hotspots (e.g., wherein the weighted sum is normalized by an estimated volume of the particular target tissue region (e.g., computed as a volume of the particular 3D volume in the function image and/or a volume of the particular target VOI))].

[0109] In certain embodiments, the particular target tissue region is selected from the group consisting of: a skeletal region comprising one or more bones of the subject, a lymph region, and a prostate region.

[0110] In certain embodiments, the 3D anatomical image is an x-ray computed tomography (CT) image, and the 3D functional image is a 3D single photon emission computed tomography (SPECT) image.

[0111] In certain embodiments, the 3D anatomical image is an x-ray computed tomography (CT) image, and the 3D functional image is a 3D positron emission tomography (PET) image.

[0112] In certain embodiments, the 3D PET image of the subject is obtained following administration to the subject of a radiopharmaceutical comprising a prostate-specific membrane antigen (PSMA) binding agent.

[0113] In certain embodiments, the radiopharmaceutical comprises [18F]DCFPyL.

[0114] In certain embodiments, at step (b), the instructions cause the processor to crop the 3D anatomical image to remove voxels representing air [e.g., to create a cropped anatomical image, and using the cropped anatomical image to identify the one or more target VOIs (e.g., using the cropped anatomical image is used as input to one or more machine learning modules, as opposed to the original size 3D anatomical image)].

[0115] In certain embodiments, the target tissue regions comprise one or more members selected from the group consisting of: a left hip bone, a right hip bone, a sacrum and coccyx region, a left clavicle, a right clavicle, a left rib, a right rib, a left scapula, a right scapula, a sternum, a lumbar vertebra, a thoracic vertebra, a skull, a cervical vertebra, a left femur, a right femur, a left humerus, a right humerus, a prostate, a urinary bladder, a rectum, a left gluteus maximus, a right gluteus maximus, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), a left kidney, a right kidney, a liver, a left lung, a right lung, a spleen, a ventricle, a left adrenal gland, a right adrenal gland, a gallbladder, a brain, a pancreas, a heart, a mandible, a left bronchi, a right bronchi, a trachea, a left common iliac artery, a right common iliac artery, and a parotid gland.

[0116] In certain embodiments, one or more of the target tissue regions comprise one or more specific bones selected from the group consisting of: a left clavicle, a right clavicle, a left femur, a right femur, a left fibula, a right fibula, a left hip bone, a right hip bone, a left humerus, a right humerus, a mandible, a left patella, a right patella, left radius, a right radius, a left tibia, a right tibia, a left ulna, a right ulna, a left rib (e.g., a first left rib, a second left rib, a third left rib, a fourth left rib, a fifth left rib, a sixth left rib, a seventh left rib, an eighth left rib, a ninth left rib, a tenth left rib, an eleventh left rib, a twelfth left rib), a right rib (e.g., a first right rib, a second right rib, a third right rib, a fourth right rib, a fifth right rib, a sixth right rib, a seventh right rib, an eighth right rib, a ninth right rib, a tenth right rib, an eleventh right rib, a twelfth right rib), a sacrum and coccyx (e.g., a combined sacrum and coccyx region; e.g., a sacrum and a coccyx individually, so as to distinguish between the two), a left scapula, a right scapula, a skull, sternum, a vertebrae region [e.g., a cervical vertebrae region, comprising one or more (e.g., up to all) cervical vertebrae; e.g., a lumbar vertebrae region, comprising one or more (e.g., up to all) lumbar vertebrae; e.g., a thoracic vertebrae region, comprising one or more (e.g., up to all) thoracic vertebrae], and an individual vertebra [e.g., e.g., an

individual cervical vertebra (e.g., a first cervical vertebra, a second cervical vertebra, a third cervical vertebra, a fourth cervical vertebra, a fifth cervical vertebra, a sixth cervical vertebra, a seventh cervical vertebra); e.g., an individual lumbar vertebra (e.g., a first lumbar vertebra, a second lumbar vertebra, a third lumbar vertebra, a fourth lumbar vertebra, a fifth lumbar vertebra, a sixth lumbar vertebra); e.g., an individual thoracic vertebra (e.g., a first thoracic vertebra, a second thoracic vertebra, a third thoracic vertebra, a fourth, thoracic vertebra, a fifth thoracic vertebra, a sixth thoracic vertebra, a seventh thoracic vertebra, an eighth thoracic vertebra, a ninth thoracic vertebra, a tenth thoracic vertebra, an eleventh thoracic vertebra, an twelfth thoracic vertebra)].

[0117] In certain embodiments, one or more of the target tissue regions comprise one or more specific bones selected from the group consisting of: a left clavicle, a right clavicle, a left hip bone, a right hip bone, a left rib (e.g., a first left rib, a second left rib, a third left rib, a fourth left rib, a fifth left rib, a sixth left rib, a seventh left rib, an eighth left rib, a ninth left rib, a tenth left rib, an eleventh left rib, a twelfth left rib), a right rib (e.g., a first right rib, a second right rib, a third right rib, a fourth right rib, a fifth right rib, a sixth right rib, a seventh right rib, an eighth right rib, a ninth right rib, a tenth right rib, an eleventh right rib, a twelfth right rib), a sacrum and coccyx (e.g., a combined sacrum and coccyx region; e.g., a sacrum and a coccyx individually, so as to distinguish between the two), a left scapula, a right scapula, a sternum, a vertebrae region [e.g., a cervical vertebrae region, comprising one or more (e.g., up to all) cervical vertebrae; e.g., a lumbar vertebrae region, comprising one or more (e.g., up to all) lumbar vertebrae; e.g., a thoracic vertebrae region, comprising one or more (e.g., up to all) thoracic vertebrae], and an individual vertebra [e.g., an individual lumbar vertebra (e.g., a first lumbar vertebra, a second lumbar vertebra, a third lumbar vertebra, a fourth lumbar vertebra, a fifth lumbar vertebra); e.g., an individual thoracic vertebra (e.g., a first thoracic vertebra, a second thoracic vertebra, a third thoracic vertebra, a fourth, thoracic vertebra, a fifth thoracic vertebra, a sixth thoracic vertebra, a seventh thoracic vertebra, an eighth thoracic vertebra, a ninth thoracic vertebra, a tenth thoracic vertebra, an eleventh thoracic vertebra, an twelfth thoracic vertebra)].

[0118] In certain embodiments, one or more of the target tissue regions comprise one or more soft-tissue regions (e.g., organs) selected from the group consisting of: a left adrenal gland, a right adrenal gland, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), a brain, a left bronchi, a right bronchi, a gallbladder, a left gluteus maximus, a right gluteus maximus, a heart, a left common iliac artery, a right common iliac artery, a left kidney, a right kidney, a liver, a left lung, a right lung, a pancreas, a prostate, a rectum, a spleen, a trachea, a urinary bladder, a ventricle, and a parotid gland.

[0119] In certain embodiments, the one or more target tissue regions comprise one or more soft-tissue regions (e.g., organs) selected from the group consisting of: a gallbladder, a left kidney, a right kidney, a liver, a left lung, a right lung, a prostate, and a urinary bladder.

[0120] In certain embodiments, the target tissue regions comprise one or more bone regions, each corresponding to a specific bone [e.g., where the target tissue regions comprise at least a portion of (e.g., up to all) the bones listed in paragraph [0116] or [0117]] and one or more soft tissue regions, each corresponding to a specific soft-tissue region [e.g., where the target tissue regions comprise at least a portion of (e.g., up to all) the soft-tissue regions listed in paragraph [0118] or [0119]].

[0121] In another aspect, the invention is directed to a system for automatically processing 3D images to automatically identify cancerous lesions within a subject, the system comprising: a processor of a computing device; and a memory having instructions stored thereon, wherein the instructions, when executed by the processor, cause the processor to: (a) receive a 3D anatomical image of a subject obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within the subject; (b) automatically identify, using one or more machine learning modules (e.g., wherein at least one of the one or more machine learning modules is a Convolutional Neural Network (CNN) module), for each of a plurality of target tissue regions, a corresponding target volume of interest (VOI) within the 3D anatomical image; (c) determine a 3D segmentation map representing a plurality of 3D segmentation masks, each 3D segmentation mask representing a particular identified target VOI (e.g., automatically, digitally stitching together the plurality of 3D segmentation masks to form the 3D segmentation map); (d) receive a 3D functional image of the subject obtained using a functional imaging modality [e.g., single-photon emission computed tomography (SPECT); e.g., positron emission tomography (PET)][e.g., wherein the 3D functional image comprises a plurality of voxels, each representing a particular physical volume within the subject and having an intensity value that represents detected radiation emitted from a the particular physical volume, wherein at least a portion of the plurality of voxels of the 3D functional image represent physical volumes within one or more of the target tissue regions of the subject]; (e) identify, within the 3D functional image, one or more 3D volume(s), each corresponding to an identified target VOI, using the 3D segmentation map (e.g., by mapping 3D segmentation masks of the 3D segmentation map to the 3D functional image); and (f) automatically detect, within at least a portion of the one or more 3D volumes identified within the 3D functional image, one or more hotspots determined to represent lesions based on intensities of voxels within the 3D functional image [e.g., based on a comparison of

intensities within the 3D functional image with a threshold value (e.g., wherein the 3D functional image is a 3D PET image and the threshold is a particular standard uptake value (SUV) level)].

[0122] In certain embodiments, the instructions cause the processor to use the one or more detected hotspots (e.g., intensities of the one or more detected hotspots) to determine a cancer status for the subject.

[0123] In certain embodiments, the target tissue regions comprise one or more reference tissue regions and wherein the instructions cause the processor to: use the 3D segmentation map to identify, within the 3D functional image, one or more 3D reference volume(s), each corresponding to a particular reference tissue region; determine one or more reference intensity values, each associated with a particular 3D reference volume of the one or more 3D reference volume(s) and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity within the particular 3D reference volume; determine one or more individual hotspot intensity values, each associated with a particular hotspot of at least a portion of the detected one or more hotspots and corresponding to a measure of (e.g., an average/mean, a median, a maximum, etc.) intensity of the particular hotspot; and determine one or more individual hotspot index values using the one or more individual hotspot intensity values and the one or more reference intensity values [e.g., wherein each individual hotspot index value is associated with a particular hotspot of the at least a portion of the detected one or more hotspots and determined using (e.g., based on a comparison between) (i) the individual hotspot intensity value associated with the particular hotspot and (ii) the one or more reference intensity values].

[0124] In certain embodiments, the reference tissue regions comprise one or more members selected from the group consisting of: a liver, an aorta (e.g., a thoracic aorta portion; e.g., an abdominal aorta portion), and a parotid gland.

[0125] In certain embodiments, the instructions cause the processor to determine an overall index value indicative of a cancer status of the subject using at least a portion of the one or more hotspot index values.

[0126] In certain embodiments, the overall index value is determined as a weighted sum of at least a portion [e.g., located within a 3D volume of the functional image that corresponds to a skeletal region of the subject; e.g., located within a 3D volume of the functional image that corresponds to a lymph region of the subject] of the individual hotspot index values [e.g., wherein each hotspot index value in the sum is weighted by a measure of size (e.g., 3D volume; e.g., average diameter) of the associated hotspot].

[0127] In certain embodiments, the 3D anatomical image is an x-ray computed tomography (CT) image, and the 3D functional image is a 3D positron emission tomography (PET) image.

[0128] In certain embodiments, the 3D PET image of the subject is obtained following administration to the subject of a radiopharmaceutical comprising a prostate-specific membrane antigen (PSMA) binding agent.

[0129] In certain embodiments, the radiopharmaceutical comprises [18F]DCFPyL.

[0130] In another aspect, the invention is directed to a system for automatically processing 3D images to identify, and measure uptake of radiopharmaceutical in, cancerous lesions (e.g., metastases) within a subject having or at risk for a cancer (e.g., prostate cancer, breast cancer, lung cancer; e.g., a metastatic cancer, such as metastatic prostate cancer, metastatic breast cancer, metastatic lung cancer), the system comprising: a processor of a computing device; and a memory having instructions stored thereon, wherein the instructions, when executed by the processor, cause the processor to: (a) receive a 3D anatomical image of a subject obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within the subject; (b) automatically identify using one or more machine learning modules (e.g., wherein at least one of the machine learning modules is a Convolutional Neural Network (CNN) module), within the 3D anatomical image: a first skeletal volume comprising a graphical representation of one or more bones of the subject; a first aorta volume comprising a graphical representation of at least a portion of an aorta of the subject; and a first liver volume comprising a graphical representation of a liver of the subject; (c) determine a 3D segmentation map representing a plurality of 3D segmentation masks, including a skeletal mask representing the identified first skeletal volume, an aorta mask representing the identified first aorta volume, and a liver mask representing the identified first liver volume; (d) receive a 3D functional image of the subject obtained using a functional imaging modality [e.g., single-photon emission computed tomography (SPECT); e.g., positron emission tomography (PET)][e.g., wherein the 3D functional image comprises a plurality of voxels, each representing a particular physical volume within the subject and having an intensity value that represents detected radiation emitted from a the particular physical volume, wherein at least a portion of the plurality of voxels of the 3D functional image represent physical volumes within one or more bones, an aorta portion, and/or a liver of the subject]; (e) automatically identify, within the 3D functional image, using the 3D segmentation map (e.g., by mapping the 3D segmentation masks of the 3D segmentation map to the 3D functional image): a second skeletal volume corresponding to the first identified skeletal volume, within the 3D anatomical image; a second aorta volume corresponding to the first aorta volume, identified within the 3D anatomical image; and a second liver volume corresponding to the first liver volume, identified within the 3D anatomical image; (f) automatically detect, within the second

skeletal volume, one or more hotspots determined to represent lesions based on intensities of voxels within the second skeletal volume (e.g., the one or more hotspots corresponding to localized regions of relatively high intensity, e.g., identified based on a comparison of intensities of voxels located within the second skeletal volume with a threshold value); and (g) determine, for each of the one or more detected hotspots, an individual hotspot index (e.g., indicative of a measure of radiopharmaceutical uptake in the lesions represented by the detected hotspot) value by: determining an aorta reference intensity level based on a measure (e.g., a mean, a maximum, a median, etc.) of intensity of voxels within the second aorta volume; determining a liver reference intensity level based on a measure (e.g., a mean, a maximum, a median, etc.) of intensity of voxels within the second liver volume; and for each individual detected hotspot: determining a corresponding individual hotspot intensity level based on a measure (e.g., a mean, a maximum, a median, etc.) of intensity of voxels of the detected hotspot; and determining a corresponding individual hotspot index level from the individual hotspot intensity level, the aorta reference intensity level, and the liver reference intensity level.

[0131] In certain embodiments, the instructions cause the processor to determine an overall index value indicative of a cancer status of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots.

[0132] In certain embodiments, the subject has or is at risk for prostate cancer.

[0133] In certain embodiments, the instructions cause the processor to: at step (b), automatically identify, within the 3D anatomical image, a first prostate volume comprising a graphical representation of a prostate of the subject; at step (c), include a prostate mask representing the identified first prostate volume in the determined 3D segmentation map; at step (e), automatically identify, within the 3D functional image, a second prostate volume corresponding to the first identified prostate volume, within the 3D anatomical image, at step (f), automatically detect one or more hotspots in the second prostate volume; and determine (i) an overall bone index value indicative of a lesion content (e.g., and severity) in the one or more bones of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second skeletal volume and (ii) an overall prostate index value indicative of a lesion content (e.g., and severity) in the prostate of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second prostate volume.

[0134] In certain embodiments, the subject has or is at risk for metastatic cancer (e.g., metastatic prostate cancer, metastatic breast cancer, metastatic lung cancer, and other metastatic bone cancers).

[0135] In another aspect, the invention is directed to a system for automatically processing a set of full body 3D anatomical images of varying sizes to automatically identify, within each 3D anatomical image, a plurality of 3D volumes that correspond to particular target tissue regions, the system comprising: a processor of a computing device; and a memory having instructions stored thereon, wherein the instructions, when executed by the processor, cause the processor to: (a) receive the set of 3D anatomical images of one or more subject(s) obtained using an anatomical imaging modality [e.g., x-ray computed tomography (CT); e.g., magnetic resonance imaging (MRI); e.g., ultra-sound], wherein the 3D anatomical image comprises a graphical representation of tissue (e.g., soft-tissue and/or bone) within each of the one or more subject(s), wherein the set of 3D anatomical images has a mean x-dimension, a mean y-dimension, and a mean z-dimension, at least one of which (the mean x-, mean y-, or mean z-dimension) having a standard deviation at least 3% of the corresponding mean [e.g., at least one of which (the mean x-, mean y-, or mean z-dimension) having a relative standard deviation of at least 3 %, at least 5%, at least 10%, at least 15%, or at least 20%]; and (b) automatically determine, using a localization module that implements a CNN, within each image of the set of 3D anatomical images, at least one initial VOI corresponding to a particular anatomical region (e.g., a group of related tissue, such as a pelvic region and a spin region) that comprises one or more particular associated target tissue regions (e.g., one or more specific bones, e.g., one or more specific organs), thereby identifying at least one initial VOI for the corresponding 3D anatomical image; and (c) for each initial VOI, automatically identify, using one or more segmentation modules each implementing a CNN, for each of the one or more particular target tissue regions associated with the particular anatomical region to which the initial VOI corresponds, a corresponding target VOI (e.g., a graphical representation of the particular target tissue region).

[0136] Features of embodiments described with respect to one aspect of the invention may be applied with respect to another aspect of the invention.

BRIEF DESCRIPTION OF THE FIGURES

[0137] The foregoing and other objects, aspects, features, and advantages of the present disclosure will become more apparent and better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

[0138] FIG. 1 is a diagram showing integration of various imaging modalities with an artificial intelligence (AI) platform, according to an illustrative embodiment.

[0139] FIG. 2 is an image showing a skeleton of a subject with various specific bones identified, according to an illustrative embodiment.

[0140] FIG. 3 is a block diagram showing use of full body segmentation in various functional imaging modalities, for various applications, according to an illustrative embodiment.

[0141] FIG. 4 is a diagram showing processing flow of different informational datasets by an AI module to classify, locate, and quantify cancerous lesions, according to an illustrative embodiment.

[0142] FIG. 5A is a block flow diagram of a process for automatically processing 3D images to identify 3D volumes within the 3D images that correspond to particular target tissue regions, according to an illustrative embodiment.

[0143] FIG. 5B is a block flow diagram of a process for automatically processing 3D images to identify 3D volumes and use the 3D volumes to detect hotspots representing cancerous lesions, according to an illustrative embodiment.

[0144] FIG. 6A is a block flow diagram of a process for automatically processing CT images for whole-body segmentation, according to an illustrative embodiment.

[0145] FIG. 6B is a set of CT images showing processing results of various steps of the process shown in FIG. 6A.

[0146] FIG. 7 is a set of three images used as training data for machine learning modules used to segment CT images, according to an illustrative embodiment.

[0147] FIG. 8 is a set of two images with identified volumes of interest overlaid.

[0148] FIG. 9A is a set of 2D views of a 3D PET/CT image obtained from a subject following administration of PyL™.

[0149] FIG. 9B is the same set of 2D views shown in FIG. 9A, following background correction.

[0150] FIG. 9C is a set of 2D views of a 3D PET/CT image obtained from a subject following administration of PyL™.

[0151] FIG. 9D is the same set of 2D views shown in FIG. 9C, following background correction.

[0152] FIG. 10A is a block flow diagram showing an example process for analyzing PyL™-PET images, according to an illustrative embodiment.

[0153] FIG. 10B is a PET/CT image obtained following administration of PyL™ to a subject.

[0154] FIG. 10C is a set of two CT images showing volumes of interest identified via the machine learning approaches described herein.

[0155] FIG. 10D is a set of 3 images illustrating background correction in PET images, according to an illustrative embodiment.

- [0156] FIG. 11A is a view of a 3D CT image with volumes of interest corresponding to specific bones labeled in color, according to an illustrative embodiment.
- [0157] FIG. 11B is another view of the 3D CT image with labeled volumes of interest shown in FIG. 11A.
- [0158] FIG. 12 is a CT image with an overlaid PET image showing identified background tissue regions.
- [0159] FIG. 13A is view a CT image with an overlaid PET image showing a detected lesion.
- [0160] FIG. 13B is another view of the CT image with overlaid PET image of FIG. 13A, showing the detected lesion.
- [0161] FIG. 13C is another view of the CT image with overlaid PET image of FIG. 13A, showing the detected lesion.
- [0162] FIG. 13D is another view of the CT image with overlaid PET image of FIG. 13A, showing the detected lesion.
- [0163] FIG. 14A is a set of two views of a CT image with an overlaid PET image obtained for a subject following administration of PyLTM.
- [0164] FIG. 14B shows the same views as FIG. 14A, but following processing via the machine learning approaches described herein to provide for background correction and detection of lesions.
- [0165] FIG. 15A is a set of two views of a CT image with an overlaid PET image obtained for a subject following administration of PyLTM.
- [0166] FIG. 15B shows the same views as FIG. 15A, but following processing via the machine learning approaches described herein to provide for background correction and detection of lesions.
- [0167] FIG. 16A is a view of a 3D CT image with identified volumes of interest overlaid.
- [0168] FIG. 16B is another view of the 3D CT image with identified volumes of interest overlaid shown in FIG. 16A.
- [0169] FIG. 17 is a block diagram of an exemplary cloud computing environment, used in certain embodiments.
- [0170] FIG. 18 is a block diagram of an example computing device and an example mobile computing device used in certain embodiments.
- [0171] FIG. 19 is a set of three graphs, each graphs showing overall index values quantifying PyLTM-PSMA uptake in lesions identified in a particular tissue region based on PET/CT image analysis for six different cancer indications, according to an illustrative embodiment.
- [0172] FIG. 20 is a graph showing individual lesion PSMA index (LPI) values determined for lesions detected in three different tissue regions (bone, lymph, and prostate) via PET/CT image analysis, according to an illustrative embodiment.

[0173] FIG. 21 is a set of two graphs comparing reference values determined via manual and automated image analysis for a blood pool reference value and a liver reference value according to an illustrative embodiment.

[0174] The features and advantages of the present disclosure will become more apparent from the detailed description set forth below when taken in conjunction with the drawings, in which like reference characters identify corresponding elements throughout. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements.

DETAILED DESCRIPTION

[0175] It is contemplated that systems, architectures, devices, methods, and processes of the claimed invention encompass variations and adaptations developed using information from the embodiments described herein. Adaptation and/or modification of the systems, architectures, devices, methods, and processes described herein may be performed, as contemplated by this description.

[0176] Throughout the description, where articles, devices, systems, and architectures are described as having, including, or comprising specific components, or where processes and methods are described as having, including, or comprising specific steps, it is contemplated that, additionally, there are articles, devices, systems, and architectures of the present invention that consist essentially of, or consist of, the recited components, and that there are processes and methods according to the present invention that consist essentially of, or consist of, the recited processing steps.

[0177] It should be understood that the order of steps or order for performing certain action is immaterial so long as the invention remains operable. Moreover, two or more steps or actions may be conducted simultaneously.

[0178] The mention herein of any publication, for example, in the Background section, is not an admission that the publication serves as prior art with respect to any of the claims presented herein. The Background section is presented for purposes of clarity and is not meant as a description of prior art with respect to any claim.

[0179] Documents are incorporated herein by reference as noted. Where there is any discrepancy in the meaning of a particular term, the meaning provided in the Definition section above is controlling.

[0180] Headers are provided for the convenience of the reader – the presence and/or placement of a header is not intended to limit the scope of the subject matter described herein.

i. Definitions

[0181] As used herein, “radionuclide” refers to a moiety comprising a radioactive isotope of at least one element. Exemplary suitable radionuclides include but are not limited to those described herein. In some embodiments, a radionuclide is one used in positron emission tomography (PET). In some embodiments, a radionuclide is one used in single-photon emission computed tomography (SPECT). In some embodiments, a non-limiting list of radionuclides includes ^{99m}Tc , ^{111}In , ^{64}Cu , ^{67}Ga , ^{68}Ga , ^{186}Re , ^{188}Re , ^{153}Sm , ^{177}Lu , ^{67}Cu , ^{123}I , ^{124}I , ^{125}I , ^{126}I , ^{131}I , ^{11}C , ^{13}N , ^{15}O , ^{18}F , ^{153}Sm , ^{166}Ho , ^{177}Lu , ^{149}Pm , ^{90}Y , ^{213}Bi , ^{103}Pd , ^{109}Pd , ^{159}Gd , ^{140}La , ^{198}Au , ^{199}Au , ^{169}Yb , ^{175}Yb , ^{165}Dy , ^{166}Dy , ^{105}Rh , ^{111}Ag , ^{89}Zr , ^{225}Ac , ^{82}Rb , ^{75}Br , ^{76}Br , ^{77}Br , ^{80}Br , ^{80m}Br , ^{82}Br , ^{83}Br , ^{211}At and ^{192}Ir .

[0182] As used herein, the term “radiopharmaceutical” refers to a compound comprising a radionuclide. In certain embodiments, radiopharmaceuticals are used for diagnostic and/or therapeutic purposes. In certain embodiments, radiopharmaceuticals include small molecules that are labeled with one or more radionuclide(s), antibodies that are labeled with one or more radionuclide(s), and antigen-binding portions of antibodies that are labeled with one or more radionuclide(s).

[0183] As used herein, “3D” or “three dimensional” with reference to an “image” means conveying information about three dimensions. A 3D image may be rendered as a dataset in three dimensions and/or may be displayed as a set of two-dimensional representations, or as a three-dimensional representation.

[0184] As used herein, an “image” – for example, a 3-D image of a subject – includes any visual representation, such as a photo, a video frame, streaming video, as well as any electronic, digital or mathematical analogue of a photo, video frame, or streaming video. Any apparatus described herein, in certain embodiments, includes a display for displaying an image or any other result produced by the processor. Any method described herein, in certain embodiments, includes a step of displaying an image or any other result produced via the method.

[0185] As used herein, a “subject” means a human or other mammal (e.g., rodent (mouse, rat, hamster), pig, cat, dog, horse, primate, rabbit, and the like).

[0186] As used herein, “administering” an agent means introducing a substance (e.g., an imaging agent) into a subject. In general, any route of administration may be utilized including, for example, parenteral (e.g., intravenous), oral, topical, subcutaneous, peritoneal, intraarterial, inhalation, vaginal, rectal, nasal, introduction into the cerebrospinal fluid, or instillation into body compartments.

[0187] As used herein, the terms “filter”, and “filtering”, as in a “filtering function” or a “filter”, refer to a function that operates on localized portions of an input array (e.g., a multi-

dimensional array) of data (e.g., image data, e.g., values computed by a layer of a CNN), referred to herein as “subpatches”, computing, for a given subpatch, a response value. In general, a filter is applied in a sliding window fashion across the array to compute a plurality of response values for the array. In particular, for a given multidimensional array, a subpatch of the array can be a rectangular region of the array having a specific size (e.g., having the same number of dimensions as the array). For example, for a $6 \times 3 \times 3$ array, a given $3 \times 3 \times 3$ subpatch refers to a given $3 \times 3 \times 3$ set of adjacent values (e.g., a neighborhood) of the array, such that there are five distinct $3 \times 3 \times 3$ subpatches in the $6 \times 3 \times 3$ array (each patch shifted one position over along the first dimension).

[0188] For example, a filtering function can compute, for a given subpatch of an array, a response value using the values of the subpatch. A filtering function can be applied in a sliding window fashion across an array, computing, for each of a plurality of subpatches of the array, a response value. The computed response values can be stored in an output array such that the positional correspondence between response values and the subpatches of the input array is maintained.

[0189] For example, at a first step, beginning with a subpatch in a corner of an input array, a filter can compute a first response value, and store the first response value in a corresponding corner of an output array. In certain embodiments, at a second step, the filter then computes a second response value for a second subpatch, shifted one position over along a specific dimension of the input array. The second response value can be stored in a corresponding position of the output array – that is, shifted one position over along a same dimension of the output array. The step of shifting position of the subpatch, computing a response value, and storing the response value in a corresponding position of the output array can be repeated for the full input array, along each dimension of the input array. In certain embodiments (e.g., a strided filtering approach), the subpatch for which the filter computes a response value is shifted more than one position at a time along a given dimension, such that response values are not computed for every possible subpatch of the input array.

[0190] As used herein, the term “convolutional neural network (CNN)” refers to a type of artificial neural network where at least one layer performs one or more filtering functions. As used herein, the term “convolution layer” refers to a layer of a CNN that receives as input an input array and computes an output array, wherein values of the output array are computed by applying one or more filters to the input array. In particular, in certain embodiments, a convolution layer receives as input an input array having $n + 1$ dimensions and produces an output array also having $n + 1$ dimensions. The first n dimensions of input and output arrays operated on by filtering layers of a CNN are referred to herein as “spatial dimensions”. The $(n + 1)^{\text{th}}$ dimension of the input is referred

to herein as the “input channel” dimension. The size of the input channel dimension is referred to herein as the “number of input channels”. The $(n + 1)^{\text{th}}$ dimension of the output is referred to herein as the “output channel” dimension. The size of the input channel dimension is referred to herein as the “number of output channels”.

[0191] In certain embodiments, a convolution layer computes response values by applying a filter that operates on subpatches that are smaller than the input array along the spatial dimensions, but extend across the full output channel dimension. For example, an $N \times M \times L \times K_0$ size input array, has three spatial dimensions and K_0 output channels. Filters of a convolution layer may operate on subpatches having sizes of $N_f \times M_f \times L_f \times K_0$, where $N_f \leq N$, $M_f \leq M$ and $L_f \leq L$. Often, a filter of a convolutional layer operates on subpatches having sizes where $N_f < N$, $M_f < M$ and/or $L_f < L$. For example, in certain embodiments, $N_f \ll N$, $M_f \ll M$ and/or $L_f \ll L$.

[0192] Accordingly, for each of one or more filters applied by a convolution layer, response values computed by a given filter are stored in a corresponding output channel. Accordingly, a convolution layer that receives an input array having $n + 1$ dimensions computes an output array also having $n + 1$ dimensions, wherein the $(n + 1)^{\text{th}}$ dimension represents the output channels corresponding to the one or more filters applied by the convolution layer. In this manner, an output array computed by a given convolution layer can be received as input by a subsequent convolution layer.

[0193] As used herein, the term “size” in reference to a filter of a convolution layer refers to a size along spatial dimensions of subpatches on which the filter operates (e.g., the subpatch size along the output channel dimension is taken as the full number of output channels). As used herein, the term “size”, in reference to a convolution layer, as in “size of a convolution layer” refers to a size of filters of the convolution layer (e.g., each filter of the convolution layer having a same size). In certain embodiments, a filter of a convolution layer has a number of variable parameters that are determined via a machine learning training process. In certain embodiments, the number of parameters of a given filter equals the number of values in a subpatch that the given filter operates on. For example, a size $N_f \times M_f \times L_f$ filter that operates on an input array with K_0 output channels has $N_f \times M_f \times L_f \times K_0$ parameters. In certain embodiments, a filter is implemented as an array, and the response value determined by the filter for a given subpatch is computed as a dot product between the filter and the given subpatch.

[0194] As used herein, the term “fully convolutional neural network (FCNN)” refers to a CNN wherein each layer of the CNN is a convolution layer.

[0195] As used herein, the term “volume”, as used in reference to an input or output of a layer of a CNN refers to an input array received or an output array computed by a CNN layer.

[0196] As used herein, the term “CNN module” refers to a computer implemented process that implements a specific CNN in order to determine, for a given input, such as an image (e.g., a 2D image; e.g., a 3D image) one or more output values. For example, a CNN module may receive as input a 3D image of a subject (e.g., a CT image; e.g., an MRI), and for each voxel of the image, determine a value that represents a likelihood that the voxel lies within a region of the 3D image that corresponds to a representation of a particular organ or tissue of the subject. A CNN module may be software and/or hardware. For example, a CNN module may be implemented entirely as software, or certain functions of a CNN module may be carried out via specialized hardware (e.g., via an application specific integrated circuit (ASIC)).

[0197] As used herein, the term “tissue” refers to bone (osseous tissue) as well as soft-tissue.

[0198] In certain embodiments, the approaches described herein can be used to segment and identify target tissue regions within a full body image of a subject. As used herein, the terms “full body” and “whole body” used (interchangeably) in the context of segmentation refer to approaches that evaluate a majority (e.g., greater than 50%) of a graphical representation of a subject’s body in a 3D anatomical image to identify target tissue regions of interest. In certain embodiments, full body and whole body segmentation refers to identification of target tissue regions within at least an entire torso of a subject. In certain embodiments, portions of limbs are also included, along with a head of the subject.

A. Detecting and Assessing Cancer Status via Nuclear Medicine Imaging and Automated Image Segmentation

[0199] Described herein are systems and methods that provide artificial intelligence (AI) – based segmentation technology that provides a basis for detecting, evaluating, and making predictions about cancer status of subjects. In particular, the AI-based segmentation techniques described herein allow for 3D representations of various tissue regions in medical images to be accurately and rapidly identified. The AI-based segmentation technologies described herein utilize machine learning techniques, such as Convolutional Neural Networks (CNNs) to automatically to identify a plurality of target 3D volumes of interest (VOIs) each corresponding to a specific target tissue region, such as one or more organs, portions of organs, particular bone(s), a skeletal region etc. Each identified 3D VOI may be represented via a segmentation mask. The multiple segmentation masks, identifying multiple target tissue regions across a patient’s body, can be stitched together to form a segmentation map. The segmentation map, and/or various segmentation masks that it comprises, may be used compute various quantities from medical

images, such as useful indices that serve as measures and/or predictions of cancer status, progression, and response to treatment. Segmentation maps and masks may also be displayed, for example as a graphical representation overlaid on a medical image to guide physicians and other medical practitioners.

[0200] In certain embodiments, AI-based segmentation is performed using an anatomical image that provides detailed structural information about locations and extent of tissue within a subject's body. Examples of anatomical images include, without limitation, x-ray computed tomography (CT) images, magnetic resonance images (MRI), and ultra-sound. Image contrast in anatomical images is a function of physical properties of underlying tissue, such as density, water and fat content. As described in further detail herein, the AI-based segmentation techniques of the present disclosure analyze contrast variations and patterns in anatomical image to identify target 3D VOIs that correspond to different specific target tissue regions.

[0201] In certain embodiments, structural information and identified target VOIs from anatomical images are combined with and/or used to analyze images obtained via functional imaging modalities. Functional images reflect physiological properties and activity in a subject. They are often acquired using probes and have intensity variations that reflect a spatial distribution of the probes within an imaged portion of a subject. For example, nuclear medicine images (e.g., PET scans; e.g., SPECT scans; e.g., whole-body scans; e.g. composite PET-CT images; e.g., composite SPECT-CT images) detect radiation emitted from the radionuclides of radiopharmaceuticals to form an image. The distribution of a particular radiopharmaceutical within a patient may be determined by biological mechanisms such as blood flow or perfusion, as well as by specific enzymatic or receptor binding interactions. Different radiopharmaceuticals may be designed to take advantage of different biological mechanisms and/or particular specific enzymatic or receptor binding interactions and thus, when administered to a patient, selectively concentrate within particular types of tissue and/or regions within the patient. Greater amounts of radiation are emitted from regions within the patient that have higher concentrations of radiopharmaceutical than other regions, such that these regions appear brighter in nuclear medicine images. Accordingly, intensity variations within a nuclear medicine image can be used to map the distribution of radiopharmaceutical within the patient. This mapped distribution of radiopharmaceutical within the patient can be used to, for example, infer the presence of cancerous tissue within various regions of the patient's body..

[0202] Transferring a segmentation map to a functional image (e.g., a nuclear medicine image, such as a PET image or a SPECT image) provides valuable context for evaluating and deriving meaning from intensity fluctuations in the functional image. In particular, it allows for regions of the functional image to be identified as corresponding to particular tissue regions of the subject. In this

manner, intensity values and fluctuations in and across voxels within a particular region can thus be understood as originating from underlying accumulation of a radiopharmaceutical in a specific tissue region (e.g., organ or bone) of the subject.

[0203] As described in further detail herein, this anatomical context serves a variety of uses. For example, it allows for measures of radiopharmaceutical uptake in various specific target organs and tissue regions, such as prostate, lymph nodes, and bone, to be assessed, and used as an indicator of presence of cancerous tissue therein. Identifying a what particular region intensities of a functional image are associated with also allows them to be evaluated in proper context. For example intensities of a certain value, if associated with a prostate region, may be more likely to be indicative of cancer than intensities of the same value, when found within another organ, such as a bladder or kidney. In certain embodiments, reference regions are identified and used to compute normalization values and levels on a scale to which intensities in other regions are compared to evaluate a cancer status. Background regions, to either be excluded from the image or used to correct for artifacts may also be identified.

[0204] Accordingly, the AI-based segmentation approaches described herein can serve as a platform on which a variety of tools and techniques for detecting, evaluating, and predicting cancer status for patients can be built. Moreover, since the AI-based segmentation tools described herein can identify a variety of important tissue regions across an entire body of a patient, and then transfer those identifications (e.g., segmentation maps) from anatomical images to various functional images, they can be used as a building block for analysis techniques based on a variety of different imaging methods that obtain contrast via a variety of different radiopharmaceuticals.

[0205] FIGs. 1 – 3 illustrate this concept. In particular, FIG. 1 is a schematic illustrating how various anatomical and functional imaging modalities can be integrated in the context of a platform agnostic AI-based segmentation technology. As illustrated in FIG. 1, the AI-segmentation platform described herein can be used as a basis for analysis of a variety of images, such as anatomical CT image themselves, ^{18}NaF images, C-11 images, Ga-68 images, FDG image, and PyL^{TM} images. FIG. 2 shows results of an example segmentation performed on a skeleton of a subject, with different specific bones identified by different colors. The segmentation shown in FIG. 2, performed using a CT image, can be transferred to a functional image in order to identify 3D volumes in the functional image that correspond to the target VOIs identified within the CT image. Other VOIs, such as soft-tissue, may also be identified and used to identify corresponding 3D volumes in a functional image. For example, FIG. 3 is a block flow diagram illustrating how segmentation of various soft-tissue and bone regions based on an anatomical image (e.g., a CT image) can be used and mapped to images obtained via various functional imaging modalities using various probes (e.g., $^{99\text{m}}\text{Tc-MIP-1404}$, $^{18}\text{F-}$

PyL, ^{68}Ga -PSMA-11, ^{18}F -NaF, ^{11}C -Choline, [18F]FDG, [18F]FACBC, etc.) to perform image based diagnosis and assessment of disease progression, state, and the like. As shown in FIG. 3, the full body segmentation techniques described herein can serve as a platform that a variety of application specific image analysis techniques that utilize different functional imaging modalities and probes can leverage and be built upon.

[0206] FIG. 4 is a diagram showing how artificial intelligence-based image analysis 400 can be performed to classify and locate lesions, as well as quantify positive lesions, e.g., for assessing cancer state in a subject. As shown in FIG. 4, a machine learning module 402, such as a convolutional neural network (CNN), receives contextual anatomical information (e.g., a CT image or other anatomical imaging modality) 404 along with functional information (e.g., a functional image, such as a SPECT, PET, or other functional image obtained using a particular imaging probe) 406. The machine learning module 402 then uses the functional 406 and anatomical 404 information to identify, classify, and/or quantify cancerous lesions in a subject 410. The machine learning module 402 may also leverage additional information 408, such as subject demographic data (e.g., age, weight, etc.) and other diagnostic test results (e.g., a biopsy, Gleason score, genetic testing results, a prostate specific antigen (PSA) test result), clinical treatment(s) (e.g., external beam radiation, brachytherapy, prostatectomy, antiandrogens, immunotherapy (e.g., Provenge), chemotherapy, etc.) as input and base the classification, identification, and/or quantification of cancerous lesions on this additional information as well.

[0207] FIG. 5 is a block flow diagram showing an example process 500 for performing segmentation to identify a plurality of target tissue regions. As shown in FIG. 5, a machine learning module receives a 3D anatomical image (e.g., a full body image) as input 502, and identifies, for each of a plurality of target tissue region, a target volume comprising a graphical representation of the target tissue region within the 3D anatomical image 504. The target volumes are stitched together 506 to form a segmentation map 510 that comprises a plurality of segmentation masks 512, with each segmentation mask representing an identified target volume. The segmentation map 510 can be stored and/or provided for display and/or further processing 508.

i. Image Segmentation using Convolutional Neural Networks (CNNs)

[0208] In certain embodiments, the AI-based image segmentation approaches described herein utilize machine learning techniques, such as CNNs to perform segmentation. As described herein, one or more machine learning modules may be used to perform intermediate analysis and processing of images, and ultimately identify target VOIs.

[0209] In certain embodiments, the AI-based image segmentation described herein utilizes a two-step approach to identify one or more of the 3D target VOIs representing target tissue regions of interest in an anatomical image. In this two-step approach, an initial volume of interest (VOI) that comprises one or more of the target VOIs is first identified. The initial VOI may be identified using a first module, referred to as a localization module. The localization module may be a first machine learning module, such as a first Convolutional Neural Network (CNN) module. Once the initial VOI is identified, a second module, referred to as a segmentation module [e.g., that implements a second machine learning network (e.g., a second CNN) to perform fine segmentation], is used to identify one or more target VOIs within the initial VOI. This approach is advantageous as allows the segmentation module to operate on a standardized input, of manageable size, without regard to variations in dimensions and boundaries of the original anatomical image received. This two-step approach is especially useful for performing accurate, finely resolved segmentation of tissue regions within a large image, such as a full body scan. The localization module may perform a rough analysis, such as by performing a coarse segmentation, to identify the initial VOI. Freed from having to accommodate varying image sizes, and analyze an entire image, the segmentation module can devote resources to performing fine segmentation to accurately identify the target VOIs within the initial VOI.

[0210] Moreover, as described in further detail herein, for example in Example 1, to identify a set of target VOIs spread out across an entire body of a subject, multiple initial VOIs can be identified in order to partition a whole body image into multiple manageable initial VOIs, each comprising a subset of the desired target VOIs to be identified. The partitioning of a whole body image into multiple initial VOIs can be performed using a single localization module, such as a machine learning module that performs a coarse segmentation to identify general anatomical regions such as an upper body, a lower body, a spine, and a pelvic region. In certain embodiments, multiple localization modules (e.g., each tasked with identifying one or more initial VOIs) may be used. For each identified initial VOI, one or more segmentation modules may then be used to perform fine segmentation and identify the one or more desired target VOIs within the initial VOI.

[0211] Various localization and segmentation modules may be combined and implemented as a single module and/or a single software application, or may be implemented separately, e.g., as separate software applications.

[0212] A number of different approaches may be used by a localization module to identify a particular initial VOI. For example, in one approach, the localization module may implement a CNN that receives a grayscale CT image (e.g., a single input channel) as input and outputs coordinates of opposite corners of a rectangular VOI (e.g., a bounding box). In another approach, the localization

module may implement a CNN that receives two input channels: (i) a grayscale CT image, and (ii) a pre-processed version of the CT image. The pre-processed version of the CT image may be a thresholded version of the CT image, wherein the thresholding provides a rough identification of the initial VOI to be identified. The CNN is trained to analyze both these input channels in order to output coordinates representing opposite corners of a rectangular initial VOI (e.g., a bounding box).

[0213] In a third approach, in order to identify a particular initial VOI comprising one or more specific target VOIs, the localization module performs a coarse segmentation to roughly identify the specific target VOIs. In this way, the localization module is essentially a rough version of the segmentation modules used to identify the specific target VOIs within the particular initial VOI. A rectangular initial VOI is then generated using the coarsely identified specific target VOIs (e.g., by drawing the smallest box that fits them, or maybe adding some buffer distance). A distinction here is that the output of the localization module in this case is not merely coordinates of cuboid vertices. In this approach, likelihood values are automatically determined for each voxel of the image that give the likelihood as to how the voxel is classified – e.g., whether the voxel is a particular initial VOI, or background, for example.

[0214] In certain embodiments, voxel classifications performed by the machine learning modules may be refined via an iterative process in order to mitigate noise in the classification process. In particular, in certain embodiments, a machine learning module such as a CNN module receives an entire image, or portion thereof as input, and outputs, for each particular voxel, a label classifying the particular voxel as (i) one of a set of pre-defined classes corresponding to anatomical regions or target tissue regions, or (ii) background. For example, a localization module trained to identify initial VOIs corresponding to a spine, a pelvic region, and a left and right upper body may label each voxel as 1, 2, 3, or 4 for spine, pelvic region, left upper body, and right upper body, respectively, or 0 for background.

[0215] A segmentation module that identifies one or more target tissue regions within a particular initial VOI may receive the particular initial VOI as input, and output an integer label for each voxel, with the integer label identifying that voxel as corresponding to one of the particular target tissue regions, or background. For example, a segmentation module that identifies different vertebrae may receive as input an initial VOI corresponding to the spine, and label each voxel in the initial VOI with an integer corresponding to a particular vertebra (e.g., 1 for a first lumbar vertebra, 2 for a second lumbar vertebra, 3 for a third lumbar vertebra, etc. with other integer labels for other bones).

[0216] The classification performed by machine learning modules as described above, may produce noisy results, with certain isolated voxels having different labels than their surroundings.

For example, a CNN module may produce an output wherein a majority of voxels within a large volume are assigned a first label, with isolated voxels and/or small groups of voxels assigned different, e.g., second or third labels, identifying them as corresponding to different anatomical regions than their neighbors. Often such isolated voxels and/or islands appear on the edges of a large uniformly labeled region.

[0217] These small non-uniformities correspond to noise in the classification process, and can be removed via an iterative procedure as follows. First, for each label representing (i) a particular anatomical region or target tissue region, or (ii) background, an associated largest connected component of labeled voxels is identified. Second, for each particular label, differently labeled isolated voxels and/or voxel islands that are immediately adjacent to voxels of the largest connected component associated with the particular label are identified and re-assigned the particular label. In certain embodiments, in a third step, any remaining isolated differently labeled voxel islands can then be removed. In a fourth step, holes in segmentation regions – voxels that are labeled as background but are completely encapsulated by non-background labeled voxels are filled (e.g., by assigning voxels of the holes a label of the non-background region that encapsulates them). The second through fourth steps, in which isolated voxels and/or voxel islands are relabeled based on their surroundings are repeated until convergence – i.e., no change from one iteration to the next. This approach reduces and/or eliminates isolated voxels of a different label from their surrounding thereby mitigating noise in the classification process.

[0218] These various approaches for identifying initial VOIs may be used alone or in combination to identify multiple initial VOIs within an anatomical image of a subject. Once one or more initial VOIs are identified, one or more target VOIs within each initial VOI may be identified. For example, PCT Publication WO2019/136349, the content of which is hereby incorporated by reference in its entirety, provides further detail on how a single initial VOI corresponding to a pelvic region may be identified within a CT image, and how fine segmentation may be performed to identify target VOIs corresponding to organs such as a prostate, bladder, gluteal muscles, as well as various pelvic bones, such as a sacrum, coccyx, and left and right hip bones, may be identified. The approaches described herein, and in particular in Examples 1 – 4, show how multiple initial VOIs can be identified and used to identify a plurality of target VOIs across a subject's body, thereby providing for accurate whole-body image segmentation that can serve as a basis for detection and assessment of cancer status, progression, and response to treatment at localized stages in various organs and tissue regions, as well as metastatic stages where it is present in multiple regions in a patient's body.

ii. Providing Anatomical Context to 3D Functional Images

[0219] Segmentation maps generated by automated AI-based analysis of anatomical images can be transferred to 3D functional images in order to identify, within the 3D functional image, 3D volumes corresponding to the target VOIs identified in the anatomical image. In particular, in certain embodiments, the individual segmentation masks (of the segmentation map) are mapped from the 3D anatomical image to the 3D functional image. The 3D volumes identified within the 3D functional image can be used for a variety of purposes in analyzing images for assessment of cancer status.

[0220] Certain identified 3D volumes and corresponding target VOIs correspond to tissue regions where cancer is suspected and/or may occur. Such regions may include, for example, a prostate, breast tissue, lung(s), brain, lymph nodes, and bone. Other regions may also be evaluated. Certain regions, such as prostate, breast, and lungs, are relevant for detecting cancer at earlier, localized stages, while others, such as lymph nodes and bone are relevant for assessing metastatic cancer. Since intensities of functional images, such as PET and SPECT images map spatial distribution of radiopharmaceutical accumulation in a patient's body, by accurately identifying specific 3D volumes in functional images that correspond to specific tissue regions intensities of voxels within those 3D volumes can be used determine a measure of uptake of radiopharmaceutical probes within the specific tissue regions to which they correspond. Since radiopharmaceutical probes can be designed to selectively accumulate in cancerous tissue (e.g., via enhanced affinity to biomolecules that are overexpressed in cancerous tissue, such as prostate specific membrane antigen (PSMA)), detecting and quantifying uptake of particular probes in certain target tissue regions and organs is of interest is indicative of and/or can be used to determining a cancer status for the subject. For example, as described in PCT Publication WO2019/136349, the content of which is hereby incorporated by reference in its entirety, assessing ¹⁴⁰⁴P uptake in a prostate volume can be used to determine a prostate cancer status for a subject. Other probes may be used to assess metastatic cancer, present in a wide variety of other tissue regions, including bones. For example, Examples 2 to 6 describe segmentation used to identify accumulation of ¹⁸F-PyLTM in cancerous lesions throughout a patient's body.

[0221] Accordingly, in certain embodiments, segmentation performed on an anatomical image, such as CT image, is transferred (e.g., mapped) to a co-registered functional image, such as a PET or SPECT image, allowing for specific tissue volumes of interest within the functional image that correspond to particular tissue regions of interest to be identified. Accumulation of radiopharmaceutical within those particular tissue regions can then be quantified based on intensities of functional image voxels that lie within the specific tissue volumes of interest.

[0222] A variety of approaches can be used to analyze voxel intensities in functional images. For example, an average, a median, a total, a maximum, etc. intensity within a specific volume may be computed and used for quantification. This computed value may then be compared with other values, e.g., computed for other tissue regions (e.g., for normalization), and used to assign a cancer status to a patient (e.g., based on comparison with one or more predetermined threshold values). In certain embodiments, localized regions of high intensity – referred to as hotspots, are identified within particular 3D volumes. As described in further detail in section B, below, these localized hotspots can be identified as representing cancerous lesions, and used to determine cancer status of a subject. In certain embodiments, machine learning approaches can be used. For example, intensities of functional image voxels lying within one or more specific tissue volumes of interest can be used as inputs to machine learning modules that compute a risk index that can, in itself be used to quantify a risk of cancer, metastasis, or recurrence of cancer, etc. and/or compared with reference values, such as thresholds, to assign a particular cancer status.

[0223] In certain embodiments, in addition to identifying specific tissue regions in which cancerous tissue is may be present, e.g., in order to determine the presence and/or the state of cancer therein, other additional tissue regions may be identified. Such additional tissue regions may correspond to background regions in which a particular radiopharmaceutical probe accumulates under normal circumstances, without lesions necessarily being present.

[0224] In certain embodiments, identified background regions are used to normalize voxel intensities so as to standardize intensity values from image to image. For example, as described in PCT Publication WO2019/136349, the content of which is hereby incorporated by reference in its entirety, gluteal muscles can be identified and uptake in them can be used to normalize intensities in SPECT images obtained following administration of 1404.

[0225] Identification of certain background regions can also be used to account for high probe accumulation levels in these regions. In certain embodiments, certain background regions are identified and excluded from analysis. For example, as described in Example 3, certain probes, such as PyL™ may accumulate in certain background regions under normal circumstances, absent any cancerous lesions or tissue within those regions. Accordingly, these regions may be identified and excluded from, e.g., hotspot detection analysis. Examples of such regions include kidneys, duodenum, small intestines, spleen, liver, pancreas, stomach, adrenal gland, rectum, and testes.

[0226] In certain embodiments, identifying 3D volumes in a functional image that correspond to background regions can be used to correct for intensity bleed effects, where high accumulation of radiopharmaceutical in a particular region may produce high intensities in the functional image not only in the 3D volume that corresponds to the particular region, but also in its

neighborhood. Such intensity bleed into other regions of functional images can mask useful signal. For example, radiopharmaceutical typically accumulates in high amounts in a patient's bladder. When imaged via a functional imaging modality, this high accumulation in the bladder may, via scattering effects, contribute to intensities in 3D volumes corresponding to tissue regions outside the bladder, such as a prostate. . Accordingly, accounting for this intensity bleed or "cross-talk" can improve the accuracy with which intensities in a 3D functional image are used to measure underlying radiopharmaceutical uptake.

[0227] In certain embodiments, the 3D volume corresponding to the particular background region producing the intensity bleed is dilated. The dilated background region may be excluded from analysis, so as to avoid using intensities of regions directly adjacent to the background region to determine uptake metrics or identification of hotspots.

[0228] In certain embodiments, a suppression method which models the profile of the functional intensities may also be used to adjust the intensity levels in neighboring regions to correct for the intensity bleed. In this approach, for a particular background region that produces intensity bleed, an amount of suppression, that is, intensity bleed, to remove from a particular voxel of the functional image is dependent on a distance from that voxel to a high intensity core region within a particular 3D background volume corresponding to the particular background region. The high intensity core region may be determined as a region comprising voxels having intensities above a predetermined value, or within a specific fraction of a maximum intensity in a specific region of the functional image.

[0229] In certain embodiments, this suppression is utilized if a maximum functional image intensity within a 3D volume identified as corresponding to the particular background region is more than a specific multiplier value times a determined background intensity value. Background intensity values may be determined based on intensities of voxels of the 3D functional image corresponding to specific reference volumes corresponding to specific tissue regions, such as gluteal muscles. The suppression approach may be applied to a sub-region of the 3D functional image in the vicinity of the particular background region producing intensity bleed. For example, it may be applied to a rectangular sub-region encompassing the particular region, plus a predetermined margin.

[0230] In certain embodiments, one or more intensity bleed functions are determined to perform suppression and thereby correct intensities of voxels of the 3D functional image for bleed (e.g., cross-talk). For example, the 3D functional image may be cropped to the aforementioned rectangular sub-region encompassing the particular region plus the predetermined margin. A determined background intensity value can be subtracted from intensities of voxels within the

cropped image region. Sample intensities can then be collected to determine how intensity originating from radiopharmaceutical uptake within the particular background region decreases as one moves away from the particular background 3D volume corresponding to the particular background region. Samples beginning at an extreme top, right, left, and bottom, and then moving outwards up, right, left, and down, respectively, may be used. Other directions are also possible. The sampled intensities provide curves (e.g., sets of sampled intensity data points) showing intensity decrease from the high intensity core. Template functions, such as n -th degree polynomials can be fit to these sampled curves and used to compute intensity values to be used as correction factors in the vicinity of the particular 3D background volume.

[0231] For example, PCT Publication WO2019/136349, the content of which is hereby incorporated by reference in its entirety describes how identification of a bladder region can be used to adjust prostate voxel intensities for intensity bleed through due to high accumulation in the bladder in 1404-SPECT images. In certain embodiments, similar approaches can be used for other images obtained with other probes, which may accumulate in the bladder, and/or other regions (e.g., the liver, the kidneys, etc.).

[0232] In certain embodiments, the segmentation approaches described herein assume that the 3D anatomical image includes certain anatomical regions. For example, an embodiment of the systems and methods described herein may assume that its input anatomical images always include a pelvic region, and automatically attempts to segment the pelvic region. The systems and methods described herein may also, for other anatomical regions, only segment such regions if they are included in the anatomical image input, for example first performing a determination as to whether they are present.

[0233] Approaches for performing whole body segmentation to identify target tissue regions corresponding to bone and/or soft-tissue (e.g., organs) are described in further detail in Examples 1, 2, 3, and 4 below. Also described are approaches for using such whole body segmentation for assessment of disease state, in particular, cancer in a patient.

B. Hotspot and Lesion Detection

[0234] In certain embodiments, instead of, or in addition to, quantifying overall uptake in a particular volume within the functional image, localized hotspots (localized regions of relatively high intensity) are detected. In certain embodiments, hotspots are detected via a thresholding approach – by comparing intensities of voxels within the functional image to one or more thresholds values. Groupings of voxels with intensities above a threshold may be detected as hotspots. A single, global, threshold value may be used, or, in certain embodiments, multiple region specific thresholds may

also be used. For example, segmentation of the co-registered anatomical image can be used to set different thresholds for different tissue regions used for hotspot detection. Segmentation of the co-registered image can also be used, as described herein, to remove effects of background signal, thereby facilitating hotspot detection (e.g., if a global threshold and/or multiple regional thresholds are used).

[0235] In certain embodiments, hotspots can additionally be detected using blob detection techniques. One approach is to use the Difference of Gaussians approach, where a PET image is filtered through a combination of high and low-pass filters approximated by Gaussian kernels. The filters reduce background noise and are invariant to changes in background levels in the different regions of the image (e.g. the thorax might have significantly higher background levels due to significant uptake in liver and kidneys compared to the pelvic region). This cascaded high/low-pass filter approach would allow for hotspot extraction without the utilization of fixed thresholds, but could instead identify local changes in the PET intensities. Another approach is to employ a Laplacian of a Gaussian blob detection method. The Laplacian of a Gaussian approach is a method for detecting edges in images using Laplacian and Gaussian filter kernels. By using different sizes of kernels, edges of structures of different sizes are detected. Choosing appropriate kernel sizes allows the method to detect structures that have properties common to lesions. The described approaches may be used in a stand-alone fashion where only one of the techniques is used for all regions of interest, or they may be used simultaneously, where different methods can be employed for different tissue regions as identified by the semantic segmentation of the co-registered anatomical image in conjunction with either a single global threshold or multiple local thresholds.

[0236] FIG. 5B shows an example process 520 for using the segmentation approaches described herein to detect hotspots representing cancerous lesions. In a first step, a 3D anatomical image is received 522. A one or more target VOIs corresponding to particular target tissue regions are identified within the 3D anatomical image 524 and the identifications of the target VOIs are stitched together 526 to create a segmentation map 534 comprising a plurality of segmentation masks 536. Each segmentation mask of the segmentation map represents a particular target VOI. Next, a 3D functional image (e.g., a PET image) is received 528 and the segmentation map is transferred to the 3D functional image to identify, for each target VOI, a corresponding 3D volume in the 3D functional image 530. The identified 3D volumes are then used to detect 532, within the 3D functional image, one or more hotspots corresponding to localized regions of high intensity relative to their surroundings. In particular, one or more specific 3D volumes may correspond to specific tissue regions where cancerous lesions may form, and be analyzed to detect hotspots therein. Other 3D volumes may be used to perform background correction, or intentionally be excluded from the

hotspot detection process. As described in further detail herein, still other volumes may be relevant in that they are used as reference volumes for computing useful indices that can be assigned to hotspots in order to score them, e.g., to indicate risk / severity of the underlying lesions they represent.

[0237] Hotspots may also be classified following their initial detection, e.g., as cancerous or not, and/or assigned likelihood values representing their likelihood of being a metastases. Hotspot classification may be performed by extracting hotspot features (e.g., metrics that describe characteristics of a particular hotspot) and using the extracted hotspot features as a basis for classification, e.g., via a machine learning module.

[0238] In certain embodiments, hotspot classification can also be performed without the use of machine learning. In such embodiments, anatomical knowledge can be combined with information relating to the shape and location of a detected hotspot to classify the hotspot as either a cancerous lesion or noise. For example, if a detected hotspot is located on the edge of a rib facing the liver and the detected peak is not the global maximum in the area surrounding the hotspot, it is possible to estimate whether the hotspot is a tumor or not based on the given anatomical and spatial information. The segmentation of the co-registered anatomical image can additionally be used to create intensity models of background tissue regions known to not contain cancerous cells, but where exceptionally high functional image voxel intensities are common. Such background regions may cause intensity bleed outside of the boundaries of the background region itself and impact the neighboring regions where cancerous lesions might exist. Profiles of the intensity levels can be estimated and used to subtract the estimated additional intensity levels present in neighboring tissue regions harboring cancerous lesions to allow a more robust hotspot detection.

[0239] Hotspots representing lesions may be used to determine risk indices that provide an indication of disease presence and/or state (e.g., a cancer status, similar to a Gleason score) for a patient. For example, metrics such as number of hotspots, a total summed intensity of identified hotspots, area fraction of a particular body part or region (e.g., skeleton) occupied by hotspots, and the like, may be used themselves as, and/or in calculation of, such risk indices. In certain embodiments, regions identified via the segmentation approaches described herein may be used in computation of risk indices, for example in computing metrics such as area fractions. Examples of approaches for using identified hotspots to compute risk indices is provided herein, in Examples 5 and 7.

C. Example CNN-Based Whole Body Segmentation and Lesion Detection Approaches

[0240] The following examples demonstrate use of the AI-based segmentation and hotspot detection approaches described herein for whole-body segmentation, detecting cancerous lesions, and determining useful metrics for evaluating a cancer status of a subject.

i. Example 1 – CNN Based Whole Body Segmentation

[0241] Example 1 describes an example approach for whole body segmentation. The implementation in this example uses five neural networks to segment bones within an entire torso. A first neural network is used to roughly localize different regions of the body. The results are used to divide the body into four regions. In each of these regions, a corresponding neural network is then used to perform segmentation into distinct bones. The results from all four regions are then combined into a finished result (e.g., a final segmentation map).

[0242] This approach is related to the two-step ‘bounding box’ approach described herein, wherein a first machine learning module (e.g., a localization module) is used to roughly localize an initial volume of interest (VOI) corresponding to an anatomical region that comprises one or more particular target tissue regions (e.g., a prostate). A second machine learning module (e.g., a segmentation module) then performs a fine segmentation within the initial VOI to identify target volume(s) corresponding to the target tissue region(s) within the initial VOI. In this case, for whole body segmentation, the first machine learning module (localization module) identifies multiple initial VOIs, each corresponding to a different anatomical region. Then, for each anatomical region, a corresponding secondary machine learning module (segmentation module) identifies one or more target volumes, each representing a particular target tissue region. The machine learning modules (e.g., neural networks) may be implemented as separate modules, or certain machine learning networks may be combined. For example, each secondary, segmentation network may be implemented as a separate module, or within a single module.

[0243] Accordingly, in this example, the first module (localization) is used for identifying the anatomical regions within the CT; that is, to find volumes of interest where networks of the second module can be applied to generate the segmentation map that is used for further analysis. The networks in the second module work with a full-resolution CT image, while the localization network works in low resolution, using a sub-sampled version of the CT image.

[0244] Three example versions of CNN networks used in software implementing whole-body segmentation in accordance with the approaches. Versions 1 and 2 segment 49 bones, and version 2 segments 49 bones and 8 soft-tissue regions.

CNN-Based Segmentation Platform Example Version 1

[0245] In a first example version of a CNN network used for whole body segmentation, the first machine learning module (localization module) in this example is referred to as “coarse-seg”, and was trained to identify 49 bones in sub-sampled CT images (a sub-sampling factor of 4 along each dimension). The localization module was used to differentiate regions of the body into a pelvic region, a spine, a left upper body, and a right upper body. The four fine segmentation networks were as follows:

- “fine-seg-pelvic”: Trained to identify the left and right ilium and the sacrum and coccyx;
- “fine-seg-spine”: Trained to identify 12 thoracic vertebrae, 5 lumbar vertebrae, and the sternum;
- “fine-seg-left-upper-body”: Trained to identify 12 ribs on the left side of the body, the left scapula, and left clavicle; and
- “fine-seg-right-upper-body”: Trained to identify 12 ribs on the right side of the body, the right scapula, and right clavicle.

[0246] The input image sizes for each network were as follows:

Table 1: Input image sizes for five neural networks.

Network Name	Input image size (no. slices, no. rows, no. columns)
coarse-seg	(81, 70, 104)
fine-seg-pelvic	(93, 146, 253)
fine-seg-spine	(194, 204, 94)
fine-seg-left-upper-body	(158, 187, 144)
fine-seg-right-upper-body	(158, 191, 146)

[0247] While the coarse-seg network of the localization module received, as an input image, an 3D anatomical image representing a large physical volume and comprising a majority of the subject’s body, the actual number of voxels in its input image was lower than for the other networks due to the factor of 4 sub-sampling. The number of voxels and size of the input image to the localization module was also reduced by cropping (removing) regions of the image that did not include graphical representations of tissue, but instead represent air. Removing these voxels as a pre-processing step further reduced the size of the image input to the localization module (see e.g., first two columns of Table 4 below). Reducing the size of the image input to the localization module via sub-sampling allows the coarse-seg network to trade resolution of the image on which it operates for additional filters, which allow it to e.g., account for variability in images it receives and

perform more accurate and robust coarse segmentation to identify the different initial volumes of interest.

[0248] The number of filters and parameters used in each neural network are listed in Tables 2 – 4, below:

Table 2: Number of convolutional filters in five neural networks.

Network Name	Total Number of Filters	Number of Filters in First Layer
coarse-seg	4096 + 49 (49 classes)	32
fine-seg-pelvic	4096 + 3	32
fine-seg-spine	2048 + 18	16
fine-seg-left-upper-body	2048 + 14	16
fine-seg-right-upper-body	2048 + 14	16

Table 3: Number of parameters in five neural networks

Network Name	Total params.	No. trainable params.	No. non-trainable params.
coarse-seg	5,881,978	5,878,678	3,300
fine-seg-pelvic	5,880,276	5,877,068	3,208
fine-seg-spine	1,472,815	1,471,177	1,638
fine-seg-left-upper-body	1,472,731	1,471,101	1,630
fine-seg-right-upper-body	1,472,731	1,471,101	1,630

[0249] Table 4 below also shows a variability in input image size for the raw data and five neural networks used in this example. As shown in the table, variability in input image size for the four secondary, fine segmentation networks is less than that for the first, localization network since the identification of the initial VOIs produces a more standardized input for the secondary machine learning modules.

Table 4: Input image size for raw data and five neural networks.

	Raw data	Coarse-seg (cropped by air)	Fine-seg-pelvic	Fine-seg-spine	Fine-seg-left-upper-body	Fine-seg-right-upper-body

			(cropped by coarse)	(cropped by coarse)	(cropped by coarse)	(cropped by coarse)
rows mean px	512	390	121	151	126	129
rows std px	0	39	7	19	20	21
rows min px	512	255	97	107	72	88
rows max px	512	399	136	189	170	188
rows min mm	700	308	133	146	98	120
rows max mm	700	546	186	258	232	242
columns mean px	512	399	207	67	112	113
columns std px	0	61	14	10	13	15
columns min px	512	311	155	45	84	72
columns max px	512	512	236	130	164	215
columns min mm	700	425	212	62	115	98
columns max mm	700	700	322	178	224	294
slices mean px	342	340	74	158	110	111
slices std px	99	95	6	12	11	12
slices min px	274	274	48	125	85	85

slices max px	624	624	92	208	136	140
slices min mm	822	822	144	375	255	255
slices max mm	1872	1872	276	624	408	420

CNN-Based Segmentation Platform Example Version 2

[0250] An updated, second, example version of the CNN whole-body segmentation system as described above in this example included adjustments to input image sizes, number of convolutional filters, and parameters used in the neural networks used to perform the segmentation. Tables 5, 6, and 7, as seen below, show updated values for the various parameters used in the five neural networks. Table 5 shows updated values for the input image sizes shown in Table 1, Table 6 shows updated values for the number of convolutional filters for the neural networks shown in Table 2, and Table 7 shows updated values for the number of parameters used by the neural networks shown in Table 3.

Table 5: Updated values for input image sizes for five neural networks.

Network Name	Input image size (no. slices, no. rows, no. columns)
coarse-seg	(81, 77, 99)
fine-seg-pelvic	(92, 144, 251)
fine-seg-spine	(192, 183, 115)
fine-seg-left-upper-body	(154, 171, 140)
fine-seg-right-upper-body	(154, 170, 140)

Table 6: Updated values for number of convolutional filters in five neural networks.

Network Name	Total Number of Filters	Number of Filters in First Layer
coarse-seg	4096 + 49 (49 classes)	32
fine-seg-pelvic	3328 + 3	26
fine-seg-spine	2048 + 18	16
fine-seg-left-upper-body	2048 + 14	16
fine-seg-right-upper-body	2048 + 14	16

Table 7: Updated values for number of parameters in five neural networks

Network Name	Total params.	No. trainable params.	No. non-trainable params.
coarse-seg	5,882,423	5,879,132	3,300
fine-seg-pelvic	3,881,200	3,883,812	2,612
fine-seg-spine	1,472,815	1,471,177	1,638
fine-seg-left-upper-body	1,472,731	1,471,101	1,630
fine-seg-right-upper-body	1,472,731	1,471,101	1,630

CNN-Based Segmentation Platform Example Version 3

[0251] Another, 3rd, example version of the CNN whole body segmentation approach was used to segment soft-tissue regions as well as bones. As described herein, this 3rd version included two coarse segmentation modules, which were used in parallel, referred to herein as “coarse-seg-02” and “coarse-seg-03”.

[0252] The “coarse-seg-02” module was trained to identify 49 bones in sub-sampled CT images. The “coarse-seg-03” module was trained to identify 49 bones and the liver, in sub-sampled CT images. The “coarse-seg-02” module outperformed the “coarse-seg-03” module for localization of bones and, to take advantage of the benefits of each module, both were used in parallel, to identify initial volumes of interest (e.g., “bounding boxes”) for different fine segmentation networks. In particular, in the 3rd version, seven fine segmentation networks were used. Six out of the seven fine segmentation networks used “coarse-seg-02” for initial volume of interest identification and a seventh, “fine-seg-abdomen”, used “coarse-seg-03” for the initial volume of interest identification.

[0253] The seven fine segmentation networks for this 3rd example version of the CNN whole-body segmentation system are as follows:

- “fine-seg-abdomen”: Trained to identify the liver, left and right kidney, and gallbladder;
- “fine-seg-left-lung”: Trained to identify the left lung;
- “fine-seg-right-lung”: Trained to identify the right lung
- “fine-seg-pelvic-region-mixed”: Trained to identify the left and right ilium, the prostate, the urinary bladder, and the sacrum and coccyx;
- “fine-seg-spine-bone”: Trained to identify 12 thoracic vertebrae, 5 lumbar vertebrae, and the sternum;

- “fine-seg-left-upper-body-bone”: Trained to identify 12 ribs on the left side of the body, the left scapula, and left clavicle; and
- “fine-seg-right-upper-body-bone”: Trained to identify 12 ribs on the right side of the body, the right scapula, and right clavicle.

Tables 8, 9, and 10, below, show values for the various parameters used in seven neural networks in the 3rd version of the CNN-based segmentation system.

Table 8: Input image sizes for the seven neural networks and two localization networks.

Network Name	Input image size (Input image size (no. slices, no. rows, no. columns))
coarse-seg-02	(81, 77, 99)
coarse-seg-03	(81, 77, 99)
fine-seg-abdomen	(92, 176, 259)
fine-seg-left-lung	(154, 171, 140)
fine-seg-right-lung	(154, 171, 141)
fine-seg-left-upper-body-bone	(154, 171, 140)
fine-seg-right-upper-body-bone	(154, 170, 140)
fine-seg-pelvic-region-mixed	(92, 144, 251)
fine-seg-spine-bone	(192, 183, 115)

Table 9: Number of convolutional filters in the seven neural networks and two localization networks.

Network Name	Total Number of Filters	Number of Filters in First Layer
coarse-seg-02	2278 + 49 (49 classes)	32
coarse-seg-03	2278 + 50	32

fine-seg-abdomen	1142 + 16 (not all output classes are included in the complete segmentation platform)	16
fine-seg-left-lung	1142 + 15	16
fine-seg-right-lung	1142 + 15	16
fine-seg-left-upper-body	1142 + 14	16
fine-seg-right-upper-body	1142 + 14	16
fine-seg-pelvic-region-mixed	1852 + 5	26
fine-seg-spine-bone	1142 + 18	16

Table 10: Number of parameters in seven neural networks and two localization networks.

Network Name	Total Params	No. trainable params	No. non-trainable params
coarse-seg-02	5,882,432	5,879,132	3,300
coarse-seg-03	5,882,469	5,879,167	3,302
fine-seg-abdomen	1,473,003	1,471,369	1,634
fine-seg-left-lung	1,472,752	1,471,120	1,632
fine-seg-right-lung	1,472,752	1,471,120	1,632
fine-seg-left-upper-body	1,472,731	1,471,101	1,630
fine-seg-right-upper-body	1,472,731	1,471,101	1,630
fine-seg-pelvic-region-mixed	3,883,812	3,881,200	2,612
fine-seg-spine-bone	1,472,815	1,471,177	1,638

[0254] Accordingly, this example demonstrates how segmentation approaches described herein can be used to perform efficient whole body segmentation.

ii. Example 2: Automated Segmentation of the Skeleton in Low-Dose CT and Quantification of Metastatic Prostate Cancer in [¹⁸F]DCFPyL PET

[0255] PSMA-PET/CT hybrid imaging is a promising diagnostic platform for prostate cancer patients. While manual delineation of organs in three dimensional CT images is often needed for accurate diagnostics and treatment planning, such manual delineation is a time consuming process. To address this challenge, this example demonstrates automating the process of accurate bone segmentation in whole body CT images using deep learning approaches in accordance with the whole body segmentation technology described herein. As described in this example, the anatomical information gained via such skeletal segmentation can be used to create a fully automated lesion detection algorithm in [¹⁸F]DCFPyL (PyLTM-PSMA) PET/CT images.

[0256] A deep learning algorithm based on cascaded deep learning convolutional neural networks for semantic segmentation of 12 skeletal regions was developed. In particular, the 12 skeletal regions were the thoracic and lumbar vertebrae, sinister (left) / dexter (right) ribs, sternum, sinister (left) / dexter (right) clavicle, sinister (left) / dexter (right) scapula, sinister (left) / dexter (right) ilium, and the sacrum. A training set ($N = 90$) and validation set ($N = 22$) of pairs of low-dose CT images and manually crafted segmentation maps were used to develop the deep learning algorithm. The algorithm's performance was assessed on a test set ($N = 10$) of low-dose CT images obtained from a PyLTM-PSMA study. In the test set of images, five representative body parts: sinister (left) ilium, lumbar vertebrae, sinister (left) ribs, dexter (right) scapula, and sternum were manually segmented. These manual segmentations were used as ground truth for evaluation of the automated segmentation procedure.

[0257] The automated segmentation can be used for automated lesion detection. For example, automated lesion detection approach using a hard threshold of standard uptake value (SUV) based on PET image voxel intensities can be performed..

[0258] Sorensen-Dice scores were used to evaluate accuracy of the automated segmentation. The segmentation approach achieved a Sorensen-Dice score mean and standard deviation of 0.95 and 0.024, respectively, on the training set and a mean and standard deviation of 0.93 and 0.036, respectively, on the validation set. For the test set, mean values (with standard deviation values shown in parentheses) for each of the five regions are as follows: 0.94 (0.016) for dexter (right) clavicle, 0.90 (0.023) for sinister (left) ribs, 0.92 (0.019) for sternum, 0.94 (0.033) for lumbar vertebrae, and 0.97 (0.0033) for sinister (left) ilium. The overall mean (over all body parts) was 0.93, with a standard deviation of 0.030.

[0259] Accordingly, this example demonstrates the accuracy of a fully automated segmentation approach for 12 skeletal regions in whole body low-dose CT images and use of an automated lesion detection approach for PyL™-PSMA/CT hybrid imaging.

iii. Example 3: Automated Whole Body Segmentation for PyL™-PET Image Analysis and Lesion Detection

[0260] This example demonstrates automated segmentation of 49 bones and 27 soft-tissue regions in whole body CT images using deep learning approaches in accordance with the whole body segmentation technology described herein. This example also demonstrates how the anatomical information gained via such segmentation can be used to create a fully automated lesion detection algorithm in [¹⁸F]DCFPyL (PyL™-PSMA) PET/CT images. This example also shows how the segmentation can be used to remove background signal from PET images to facilitate observation and detection of lesions in which PyL™ has accumulated.

[0261] FIG. 6A shows a block flow diagram illustrating an embodiment of the segmentation processes described herein that is used in the PyL™-PET/CT image analysis described in this example. As in process 500, shown in FIG. 5, in process 600 an anatomical image is received. In particular, in process 600, a CT image is received 610a. In this example, process 600 is used to identify and segment the CT image into target volumes of interest corresponding to 49 specific bones and 8 soft tissue regions (in this example, example version 3 of the CNN-based segmentation approach described in Example 1 is used). In order to identify the target volumes, a coarse segmentation is performed 620a to localize a set of initial volumes of interest, or sub-regions, e.g., as described in Example 1. A fine segmentation is then performed within each of the sub-regions to identify the specific target volumes of interest that correspond to the target 49 bone and 8 soft-tissue regions 630a. Segmentation masks representing the identified target volumes can be created, and merged 640a, for example to create a 3D whole body segmentation map.

[0262] As noted in Example 1 (CNN Network Version 3), the specific regions segmented are as follows:

49 Bones:

clavicle_left
clavicle_right
hip_bone_left
hip_bone_right
rib_left_1
rib_left_10

rib_left_11
rib_left_12
rib_left_2
rib_left_3
rib_left_4
rib_left_5
rib_left_6
rib_left_7
rib_left_8
rib_left_9
rib_right_1
rib_right_10
rib_right_11
rib_right_12
rib_right_2
rib_right_3
rib_right_4
rib_right_5
rib_right_6
rib_right_7
rib_right_8
rib_right_9
sacrum_and_coccyx
scapula_left
scapula_right
sternum
vertebra_lumbar_1
vertebra_lumbar_2
vertebra_lumbar_3
vertebra_lumbar_4
vertebra_lumbar_5
vertebra_thoracic_1
vertebra_thoracic_10
vertebra_thoracic_11

vertebra_thoracic_12
vertebra_thoracic_2
vertebra_thoracic_3
vertebra_thoracic_4
vertebra_thoracic_5
vertebra_thoracic_6
vertebra_thoracic_7
vertebra_thoracic_8
vertebra_thoracic_9

8 Soft-Tissue Regions:

gallbladder
kidney_left
kidney_right
liver
lung_left
lung_right
prostate
urinary_bladder

[0263] FIG. 6B shows a series of CT images overlaid with annotations illustrating the steps in the whole-body segmentation process 600 described in this example. An example raw CT image that is received at step 610a is shown in 610b. Image 620b shows results of a coarse segmentation (different regions identified are shown as colored regions) used to identify the initial volumes of interest in the CT image. Image 630b shows “bounding boxes” identifying initial volumes of interest in which fine segmentation is performed to identify the target volumes of interest corresponding to the 49 bones and 8 soft tissue regions. Image 640 illustrates the final, merged whole body segmentation.

[0264] Turning to FIG. 7, in order to train the machine learning modules used to perform the segmentations, numerous pre-labeled sample images, such as the three images (710, 720, 730) shown in FIG. 7, were used as a training data set. The machine learning modules were fine-tuned using 1000's of manually annotated labels along with 100's of models trained on dozens of GPUs over 1000's of hours before optimal configurations for the many components of the segmentation

platform were discovered. Once trained, however, segmentation can be performed rapidly, with results such as those shown in FIG. 8 (810 and 820 are two example images) and FIG. 16A and 16B obtained in a matter of seconds -typically under 180 seconds.

[0265] The anatomical context obtained by segmenting the CT image can be used to detect lesions in PET images that are co-registered with the CT image, for example, as in a typical PET/CT imaging modality in which CT and PET images are obtained for a subject in quick succession, with the subject remaining in a substantially same position as the images are recorded. In particular, as demonstrated in this example, segmentation of the CT image can be transferred to the PET image in order to identify regions of interest such as the prostate gland or the skeleton, where the entire functional image except for regions expected to carry either primary or secondary prostate cancer tumors can be excluded from the lesion detection algorithm. Without wishing to be bound to any particular theory, exclusion of regions from the lesion detection algorithm can be especially important for background regions where accumulation leads to high intensity in PET image voxels within and around the background tissue regions. The high intensities in background tissue regions may, if not excluded from the prostatic lesion detection process, lead to erroneous classification of background noise as cancerous lesions. In the regions that remain after the background exclusion process, a simple threshold (e.g., a SUV of 3) in conjunction with a lesion classification algorithm can be employed to find hotspots within the relevant tissue regions. The classification algorithm may be used as a simple check to confirm the hotspots position and to compare the intensity in the hotspot's neighborhood. If the hotspot is part of a larger hotspot and is located on the edge of a body part (e.g. a rib close to the liver), the lesion may be classified as noise and excluded. FIGs. 9A – D illustrate the need for, and benefits obtained by, this background subtraction approach and lesion detection algorithm, where only regions of interest are included in the detection procedure, and portions of the image contains high intensity background voxels. FIGs. 9A and 9C show PET images as white-to-red-to-blue false color overlaid on CT images. Blue color indicates low PET image intensity, transitioning from red to white as intensity increases. FIGs. 9B and 9D show the same PET images as in 9A and 9C, but with all background uptake removed from the PET image. In FIG. 9B and 9D, a few localized hotspots, corresponding to cancerous lesions are visible. In the PET images without background removal, these hotspots were overwhelmed and masked by the large background signal.

[0266] FIG. 10A shows an example process 1000 for using the anatomical context provided by whole body CT segmentation for detecting lesions and determining useful risk indices that quantify a cancer status (e.g., prostate cancer status) of a subject using PET/CT images obtained following administration of Pyl^{TM} as a radiopharmaceutical. As shown in FIG. 10A, PET/CT images

are obtained 1010, anatomical context is determined 1012 via segmentation of bones and soft tissue and used to obtain measurements (e.g., of lesions and/or risk indices) from the PET image 1014. FIG. 10B shows an example PET/CT image 1016 and FIG. 10C shows example segmentations of bones and soft tissue 1018a and 1018b. As shown in FIG. 10D, segmentation of bone and soft-tissue (e.g., shown in image 1020a) can be used to remove background signal (e.g., from background tissue regions, such as those shown in image 1020b), leaving only desired signal from hotspots indicative of cancerous lesions 1020c.

[0267] FIGs. 11 – 13 illustrate this process, showing how segmentation of a CT image can be transferred to a PET image, and are used to provide anatomical context for background removal and lesion detection. In particular, FIG. 11A and FIG. 11B show different views of a CT image, in which segmentation of bones has been performed. The various segmented bones are labeled with different colors in the views shown in FIG. 11A and FIG. 11B. Notably, by virtue of the “bounding box” approach described herein, even such a complex and large scale segmentation can be performed extremely rapidly. The segmentation shown in FIGs. 11A and 11B was performed in less than 180 seconds.

[0268] As shown in FIG. 12, segmentation of bones and soft tissue can be used to account for normal uptake in certain background tissue regions, which may occur under normal circumstances and mask useful hotspot signals that indicate presence of cancerous lesions. In PET/CT imaging performed with PyL™, normal uptake can occur in kidneys, duodenum, small intestines, spleen, liver, pancreas, stomach, adrenal gland, rectum, and testes. FIG. 12 shows the aforementioned regions identified in red. As shown in FIG. 13A, once background intensities due to normal PyL™ accumulation in these regions is subtracted out, lesions become readily observable and can be detected. One such lesion 1302 is visible in FIGs. 13A - D. Lesions such as lesion 1302 can be detected, for example via a thresholding approach, and classified as PyL™ positive (i.e., indicative of cancer) e.g., as described above. As with segmentation, lesion detection is also rapid. Lesion 1302 was detected in less than 5 seconds.

[0269] FIG. 14A and FIG. 14B show results of the PET image processing and lesion detection approach described above performed on PET/CT images obtained for a 65 year old male with metastatic castration resistant prostate cancer. The patient’s prostate specific antigen (PSA) score was 6.8 ng/ml. The PET images were obtained following administration of PyL™ as the radiopharmaceutical. FIG. 14A shows the initial, raw, PET/CT image. FIG. 14B shows the image following the processing described herein. The background intensity is removed, and several small lesions are readily observable. The segmentation of the 49 bones and 8 soft-tissue regions (e.g., organs) was performed in less than 180 seconds.

[0270] FIG. 15A and FIG. 15B show results of the PET image processing and lesion detection approach performed on PET/CT images obtained for a 54 year old male with metastatic castration resistant prostate cancer. The patient's PSA score was 33.55 ng/ml. The PET images were obtained following administration of PyL™ as the radiopharmaceutical. FIG. 15A shows the initial, raw, PET/CT image. FIG. 15B shows the image following the processing described herein. The background intensity is removed, and several small lesions are readily observable. The segmentation of the 49 bones and 8 soft-tissue regions (e.g., organs) was again performed in less than or about 180 seconds.

[0271] Accordingly, this example demonstrates how the whole-body segmentation approach described herein contextualizing the PyL™ images, with automated segmentation of 27 soft-tissue organs and 49 bones, to detect, quantify and track PyL™ avid lesions. The approach would allow the clinicians/physicians to ask clinically relevant questions for better management of prostate cancer patients. Advantages such as increased diagnostic accuracy, precision, speed, and reproducibility, were demonstrated (statistically) in the context of artificial intelligence assisted 1404-SPECT image analysis (see, e.g., PCT Publication WO2019/136349, the content of which is hereby incorporated by reference in its entirety), may also be obtained for PyL™-PET images.

iv. Example 4: Example Bone and Soft-Tissue Segmentation Regions

[0272] This example provides a listing of a set of example bone and soft-tissue regions that a system developed using the approaches and embodiments described herein can identify via segmentation of CT images. In particular, listed below are 67 bones and 22 soft-tissue regions for that have been manually labeled (e.g., identified by one or more human experts) in a set of CT images. These manually labeled CT images can be used as training data for the machine learning approaches described herein. For example, while Examples 1 – 3 describe current versions of software implementing whole body segmentation approaches that segment 49 bones and 8 soft-tissue regions, their functionality can readily be updated to segment any number of the 67 bone and 22 soft-tissue regions listed in this example. Accordingly, this example shows that embodiments of the systems and methods described herein may be developed to identify similar regions, including, but not necessarily limited to the specific regions described in this example. Certain systems and methods may identify tissue regions not necessarily listed in this example. In certain embodiments, both bone and soft-tissue regions are identified. In certain embodiments, some systems and methods may identify only bones, or only soft-tissue.

[0273] As indicated in the listing below, certain left and right hand side bones are identified as separate tissue regions (e.g., a left clavicle and a right clavicle) and, in certain cases, individual

members of large groups of bones are identified separately. For example, the example listing below shows that individual ribs and vertebrae are identified via the segmentation approach of this example (specific ribs and vertebrae are numbered in the listing). This example also should make clear that the approaches described herein can be used to segment a variety of regions throughout the body, including, but not necessarily limited to the regions listed herein.

Segmentation regions:

Bones (67)

clavicle_left
clavicle_right
femur_left
femur_right
fibula_left
fibula_right
hip_bone_left
hip_bone_right
humerus_left
humerus_right
mandible
patella_left
patella_right
radius_left
radius_right
rib_left_1
rib_left_2
rib_left_3
rib_left_4
rib_left_5
rib_left_6
rib_left_7
rib_left_8
rib_left_9
rib_left_10
rib_left_11
rib_left_12

rib_right_1
rib_right_2
rib_right_3
rib_right_4
rib_right_5
rib_right_6
rib_right_7
rib_right_8
rib_right_9
rib_right_10
rib_right_11
rib_right_12
sacrum_and_coccyx
scapula_left
scapula_right
skull
sternum
tibia_left
tibia_right
ulna_left
ulna_right
vertebra_cervical_all
vertebra_lumbar_1
vertebra_lumbar_2
vertebra_lumbar_3
vertebra_lumbar_4
vertebra_lumbar_5
vertebra_lumbar_6
vertebra_thoracic_1
vertebra_thoracic_2
vertebra_thoracic_3
vertebra_thoracic_4
vertebra_thoracic_5
vertebra_thoracic_6

vertebra_thoracic_7
vertebra_thoracic_8
vertebra_thoracic_9
vertebra_thoracic_10
vertebra_thoracic_11
vertebra_thoracic_12

Soft tissue (22):

adrenal_gland_left
adrenal_gland_right
aorta_abdominal_part
aorta_thoracic_part
brain
bronchi_left
bronchi_right
gallbladder
gluteus_maximus_left
gluteus_maximus_right
heart
kidney_left
kidney_right
liver
lung_left
lung_right
pancreas
prostate
rectum
spleen
urinary_bladder
ventricle

v. Example 5: Computing Hotspot Indices for Radiopharmaceutical Uptake Quantification and Clinical Endpoint Assessment

[0274] Example 5 is an example approach that uses the segmentation and hotspot detection methods described herein to compute, for a particular detected hotspot, a hotspot index value that can be used to infer and/or quantify uptake of radiopharmaceutical within the lesion that the detected hotspot represents. The computed hotspot index can be related to clinical endpoints, including survival rate of a patient and to determine treatment strategy. When computed for multiple images, collected at different time points, the computed index values can be compared with each other for a particular patient, and the change in index used to evaluate efficacy of a treatment and to make a prognosis of how the index will change in the near future. In certain embodiments, the computed index can predict sensitivity towards treatment targeting the imaging-ligand. In certain embodiments, the computed index can also be included in nomograms for effective patient stratification.

[0275] The approach in this example uses a CT-PET image set obtained for a particular patient following administration of a radiopharmaceutical comprising a PSMA binding agent, for example, PyL™. However, the approaches described herein are agnostic to the particular radiopharmaceutical used for imaging, and can be utilized with a variety of different radiopharmaceuticals, e.g., ^{99m}Tc-MIP-1404, ¹⁸F-PyL, ⁶⁸Ga-PSMA-11, ¹⁸F-NaF, ¹¹C-Choline, [18F]FDG, [18F]FACBC, etc.

[0276] A machine learning-based approach in accordance with the systems and methods described herein is used to identify various target VOIs within the CT image. As described herein, the target VOIs are volumes in the CT image identified, automatically, by the segmentation approach, to correspond particular target tissue regions. In this example approach, target VOIs corresponding to a liver, an aorta portion, and a parotid gland are identified. As described herein, these particular tissue regions – the liver, aorta portion, and parotid gland – serve as reference regions for computation of the hotspot index. Other target VOIs corresponding to other target tissue regions, in addition to the reference regions, may also be identified. Segmentation masks representing the identified target VOIs are mapped to the PET image to identify corresponding 3D volumes within the PET image. In this manner, 3D reference volumes corresponding to the liver, aorta portion, and parotid gland are identified in the PET image. Once identified, each particular 3D reference volume is used to compute a corresponding reference intensity value that provides a measure of the intensities of voxels within the particular 3D reference volume. In this example, a mean intensity inside each volume is used, though other measures (e.g., a median, a maximum, a mode, etc.) are possible.

[0277] To compute an index value for a particular identified hotspot, a hotspot intensity value is computed for that particular hotspot and compared with the reference intensity values. Similar to the reference values, the hotspot intensity value provides a measure of intensity of voxels of the hotspot. In this example, a maximum value is computed, although, as with the reference values, other measures can be used. As this example shows, the particular measure used to compute the hotspot intensity value need not be the same as that used to compute the reference intensity values. To compute the hotspot index value, the reference intensity values can be mapped to reference index values on a scale, and the hotspot index value can then be computed based on whether the hotspot intensity value lies above, below, or in between the reference values.

[0278] For example, typically the reference intensity value computed from the reference volume corresponding to the aorta portion (this – aorta – region is used to provide a measure of uptake in the blood pool, and may also be referred to as a blood or blood pool reference value) is lowest in value, followed by the liver and then the parotid region. Accordingly, a hotspot whose hotspot intensity value is equal to the blood reference value will be assigned a hotspot index value equal to 100; a hotspot whose hotspot intensity value is equal to the liver reference value will be assigned a hotspot index value equal to 200; and a hotspot whose hotspot intensity value is equal to the parotid gland reference value will be assigned a hotspot index value equal to 300. Hotspot index values for hotspot intensity values lying in between two reference intensity values can be determined via interpolation (e.g., linear interpolation).

[0279] In this manner, a detected hotspot can be assigned an index value that quantifies, in a standardized fashion (e.g., so as to be comparable between different images) a level of radiopharmaceutical uptake in the particular lesion that the hotspot represents. As described herein, these indices can be related to survival, which, in turn, makes them useful for treatment management. For example, depending on the expected outcome of the patient, a more or less aggressive treatment may be considered. Price of treatments can also be included as a factor. The index is especially useful when measured over time. When comparing indices across time and multiple imaging examinations for a patient, the change in index can be used to evaluate the efficacy of a treatment, and to make a prognosis how the index will change in the near future.

[0280] Notably, a significant advantage of the approach described in this example over previous approaches is the ability to compute reference intensity values (as well as hotspot intensity values) from automatically identified 3D volumes provided by the artificial intelligence based segmentation approaches described herein. Previous approaches that attempted to quantify regions of images representing lesions relied on hand marking – e.g., via placement of a circular marker – of regions in 2D slices to identify 2D regions of interest lying within a reference tissue

region. In contrast to such small 2D regions, the 3D volume that are identified via the approaches used herein capture intensities throughout entire organs, and thereby offer increased accuracy and repeatability. Moreover, by using accurate automated segmentation, further increases in accuracy and repeatability are provided.

[0281] In addition, rather than classifying a detected lesion using one of a small number of values, the approach of this example uses a continuously varying index computed via interpolation. This approach provides more detailed information that can be utilized to manage treatment strategy and track disease progression and/or treatment efficacy over time in a finely grained and accurate fashion.

vi. Example 6: Automated Hotspot Detection and Uptake Quantification in Bone and Local Lymph

[0282] PyLTM-PSMA PET/CT hybrid imaging (*e.g.*, images PET/CT images acquired for a patient after administering PyLTM to the patient) is a promising tool for detection of metastatic prostate cancer. Image segmentation, hotspot detection, and quantification technologies of the present disclosure can be used as a basis for providing automated quantitative assessment of abnormal PyLTM-PSMA uptake in bone and local lymph (*i.e.*, lymph nodes localized within and/or in substantial proximity to a pelvic region of a patient). In particular, as shown in this example, image segmentation and hotspot detection techniques in accordance with the systems and methods described herein can be used to automatically analyze in PET images to detect hotspots that a physician might identify as malignant lesions.

[0283] In this example, PET/CT scans were evaluated to automatically identify, within the PET images of the scans, hotspots corresponding to potential bone and local lymph node lesions. For each scan, a semantic segmentation of the CT image was performed using the deep learning approaches described herein in order to identify a set of specific bone and soft-tissue regions (*e.g.*, organs). Once obtained, the CT image segmentation was transferred (*e.g.*, mapped) to the PET image of the PET/CT scan to identify corresponding 3D volumes in the PET image. In each PET image, intensities corresponding to background uptake were removed by suppressing intensities in identified volumes corresponding to a urinary bladder, a liver, and a kidney. The segmentation was also used to define relevant volumes in which to detect hotspots. Blob detection algorithms were then applied to identify abnormal hotspots representing either possible bone lesions or possible malignant local lymph nodes. Reference volumes corresponding to a liver and a thoracic part of an aorta were used to compute reference SUV values.

[0284] Accuracy of the hotspot detection approach in this example was validated by comparing detection of hotspots in PET images using the automated approach described herein with manual annotations identifying bone lesions and lymph node lesions from teams of physicians. A set of 157 PET/CT scans that were annotated for bone lesions (114 of the images did not have any lesions and 11 of the images had greater than three lesions) were used to evaluate accuracy in detecting hotspots corresponding to bone lesions and a set of 66 scans that were annotated for local lymph node lesions (40 images without lesions and six with greater than three lesions). The bone detection algorithm identified 97% of all annotated bone lesions, with on average 109 hotspots per image. The local lymph detection algorithm found 96% of all annotated malignant local lymph nodes, with on average 32 hotspots per scan.

[0285] Accuracy of the deep learning segmentation was also evaluated in this example. The segmentation algorithm was trained to segment a set of 52 bones and 7 soft tissue regions (*e.g.*, organs) used either for defining the hotspot search region or as reference regions. Training and validation was performed and evaluated for each particular region (bone or soft tissue region) using a manual identification of that region in a CT image. For example, 140 manually identified livers were used to train the algorithm for liver segmentation, and 37 manually identified livers were used for validation. Similarly, 61 and 14 manually identified aortas were used for training and validation of the aorta region segmentation, respectively. For liver segmentation, Dice scores of 0.99 and 0.96 were obtained for the training and validation sets, respectively. For aorta segmentation, Dice scores of 0.96 and 0.89 were obtained for training and validation, respectively. Finally, a set of ten additional images that were not used development of the algorithm were used to assess generalization. For these ten images, Dice scores characterizing segmentation accuracy of the liver and aorta were 0.97 ± 0.01 and 0.91 ± 0.5 , respectively.

[0286] A full listing of the specific 52 bone (the sacrum and coccyx are listed on a single line as “sacrum_and_coccyx”, but correspond to two segmented regions) and 7 soft-tissue regions used in this particular example is below.

Bone Regions:

capula_left
clavicle_left
clavicle_left
clavicle_right
hip_bone_left
hip_bone_right
rib_left_1

rib_left_10
rib_left_11
rib_left_12
rib_left_2
rib_left_3
rib_left_4
rib_left_5
rib_left_6
rib_left_7
rib_left_8
rib_left_9
rib_right_1
rib_right_10
rib_right_11
rib_right_12
rib_right_2
rib_right_3
rib_right_4
rib_right_5
rib_right_6
rib_right_7
rib_right_8
rib_right_9
sacrum_and_coccyx
scapula_left
scapula_right
sternum
vertebra_lumbar_1
vertebra_lumbar_2
vertebra_lumbar_3
vertebra_lumbar_4
vertebra_lumbar_5
vertebra_thoracic_1

vertebra_thoracic_10
vertebra_thoracic_11
vertebra_thoracic_12
vertebra_thoracic_2
vertebra_thoracic_3
vertebra_thoracic_4
vertebra_thoracic_5
vertebra_thoracic_6
vertebra_thoracic_7
vertebra_thoracic_8
vertebra_thoracic_9

Soft Tissue Regions:

aorta_abdominal_part
aorta_thoracic_part
kidney_left
kidney_right
liver
prostate
urinary_bladder

[0287] Accordingly, this example demonstrates use of deep learning based semantic segmentation for automated identification of hotspots and for computation of SUV reference values in [^{18}F] DCFPyL (PyLTM-PSMA) PET/CT images.

vii. Example 7: Lesion PSMA Score and PSMA Weighted Lesion Involvement

[0288] This example provides an approach for assigning detected hotspots hotspot indices, based on a comparison of individual hotspot intensities with reference levels, and then using the assigned individual hotspot indices to determine overall indices representing a weighted sum of measures of lesion size. The indices used in this example are used to assess cancer status for patients imaged using a PSMA binding agent and PET-CT imaging.

[0289] In particular, individual hotspot indices are determined via interpolation from our reference levels similar to the approach described in Example 5. In this example, reference VOIs

corresponding to an aorta portion and a liver portion are segmented and mapped to corresponding 3D volumes in a PET image. The aorta portion volume is used to determine a blood reference intensity, and the liver volume is used to determine a liver reference intensity. In this example, each reference intensity is determined from its corresponding volume by taking an average intensity (SUV) in the corresponding volume, but other measures, such as a maximum, peak, or median value could also be used. Reference levels on a scale, referred to as a Lesion PSMA Score (LPS), are assigned to intensity values based on the blood and liver reference intensities (SUVs) as follows: a LPS of 0 is assigned to a 0 SUV level, an LPS of 1 is assigned to the blood reference intensity, an LPS of 2 is assigned to the liver reference intensity, and a maximum LPS of 3 is assigned to a reference intensity calculated as twice the liver reference intensity.

[0290] Individual hotspots are assigned LPS scores based on their individual intensities. For individual hotspots having intensities ranging from 0 to the maximum reference intensity (twice the liver reference intensity), the LPS score corresponding to the individual hotspot intensity is interpolated from the reference scale. Hotspots having intensities greater than the maximum reference intensity are assigned the maximum LPS of 3.

[0291] Two example overall risk indices that can be computed using detected hotspots and the individual hotspot indices (LPS scores) are calculated as weighted sums of hotspot sizes in particular volumes corresponding to tissue regions where cancerous lesions may occur. A first example index is a PSMA-weighted total bone/lymph/prostate lesion volume or ratio (PTLV or PTLR). This index is a weighted sum of lesion volumes, the weight being the lesion PSMA score. The sum is computed separately for 3D volumes corresponding to bone (e.g., a skeletal region), lymph nodes, and prostate as the weighted sum of hotspot volumes for hotspots found in each particular region. In particular, the weighted sum is computed as follows:

$$\sum (lesion\ volume \times lesion\ psma\ score)$$

[0292] In certain case, a ratio may be preferable and can be computed by dividing the weighted sum by the total 3D volume of the particular region (e.g., in the PET image). Weighting summed hotspot volumes by LPS, as opposed to, for example SUV or normalized SUV is advantageous since it PSMA expression in the form of (normalized) SUV values may not relate to aggressiveness of the disease in a linear fashion. That is, for example, it is not a given that a hotspot with an intensity of 100 represents a lesion is five times worse than one represented by a hotspot having an intensity of 20. Calculating the LPS score and weighting hotspots by the LPS score provides a scale to compare different hotspots.

[0293] Another example index is a PSMA-weighted bone/lymph aggregated diameter (PLAD). This index is also a sum of a measure of lesion size, weighted by LPS score, but instead of

volume this index uses an average diameter (e.g., averaged of x, y, and z-diameters) of each hotspot. Since volume is a three-dimensional quantity, a minor change in volume for a large lesion can dominate (e.g., cause large fluctuations in the sum) over changes in size of smaller lesions. Using the average diameter instead mitigates this effect. This index is calculated as follows:

$$\sum (lesion\ average\ diameter \times lesion\ psma\ score)$$

[0294] The weighed aggregated diameter can be calculated for the bone and lymph.

vii. Example 8: Improved Performance in AI-Assisted Image Analysis for Patients with Low or Intermediate Risk Prostate Cancer

[0295] 99mTc MIP-1404 (1404) is a PSMA targeted imaging agent for the detection and staging of clinically significant prostate cancer. Manual assessment of tracer uptake in SPECT/CT images introduces inherent limitations in inter- and intra-reader standardization. This example describes a study that evaluated the performance of PSMA-AI assisted reads, wherein automated segmentation of prostate volumes and other target tissue regions are performed in accordance with the embodiments described herein, over manual assessment and known clinical predictors.

[0296] The study analyzed 464 evaluable patients with very low-, low-, or intermediate- risk prostate cancer, whose diagnostic biopsy indicated a Gleason grade of $\leq 3+4$ and/or who were candidates for active surveillance (1404-3301). All subjects received an IV injection of 1404 and SPECT/CT imaging was performed 3-6 hours postdose. Three independent readers evaluated the images. All subjects underwent either voluntary RP (low- and intermediate-risk) or prostate biopsy (very low-risk) post dosing. Clinically significant disease was declared in subjects with Gleason grade 7 or higher. The PSMA-AI was developed and locked prior to the analysis. Three different independent readers used PSMA-AI to obtain quantitative expression of 1404 in the prostate against the background (PSMA-Index). PSMA-Index for all readers and subjects were compared to the histopathological reference, yielding 6 receiver operating characteristic (ROC) curves (3 manual reads + 3 PSMA-AI assisted reads). The clinical performance of the 1404 PSMA-AI assisted read was also evaluated by comparing the Area Under the ROC Curve (AUC), of a multivariate model (PSA, clinical staging and diagnostic Gleason score) with and without PSMA-Index.

[0297] The manual reads demonstrated AUCs of 0.62, 0.62 and 0.63. The reads with PSMA-AI demonstrated AUCs of 0.65, 0.66 and 0.66. The PSMA-AI performance in terms of AUC was higher than manual in all 3*3=9 pairwise comparisons between the two reader groups, with statistically significant improvement observed in five cases (nominal $p < 0.05$), not accounting for multiple comparisons. The predictive ability of the baseline multivariate model, without PSMA-Index, was at AUC 0.74. Upon adding of PSMA-Index, the model predictive ability increased to AUC 0.77. The

logistic regression model indicated that PSMA-Index ($p=0.004$), pre-surgery PSA (0.018) and % positive cores ($p<0.001$) were significantly associated with clinically significant disease. When measuring reproducibility, log (PSMA-Index) correlation coefficients for pairs of PSMAAI readers were 0.94, 0.97 and 0.98.

[0298] Accordingly, the study described in this example demonstrated that PSMA-AI provides a standardized platform to generate reproducible quantitative assessment of 1404. The PSMA-AI assisted read demonstrated an additive improvement over known predictors for identifying men with clinically significant disease.

ix. Example 9: Automated Calculation of PSMA Indices for Quantification of 18F-DCFPyL Uptake from PET/CT Images for Prostate Cancer Staging

[0299] This example demonstrates automated image segmentation, lesion detection, and calculation of a standardized index score based on hotspot indices assigned to detected lesions via embodiments of the approaches described herein. The automated image analysis procedures are used for evaluated cancer status of patients imaged via PET/CT scanning after being administered the radiopharmaceutical 18F-DCFPyL (PyL™).

[0300] In this example, a cascaded deep learning pipeline was used to segment relevant organs in the CT image, and segmentations are projected into PET image space. In particular, target regions corresponding to a bone volume corresponding to bones of the subject, a lymph volume corresponding to lymph regions, and a prostate volume corresponding to a prostate gland were segmented in the CT image and mapped to identify corresponding 3D volumes in the PET image. Likewise, aorta and liver volumes corresponding to an aorta portion and a liver were also segmented and mapped to the PET image for use as reference regions as described herein.

[0301] Hotspots were detected in the PET image. Hotspots that were considered to represent lesions are manually segmented, as well as additional lesions not detected by the algorithm. Each individual detected hotspot was quantified by calculating an individual hotspot index referred to as lesion miPSMA index (LPI). As described in Examples 6 and 7 above, the LPI used in this example is a continuous index computed using automatically segmented (as opposed to manually identified) volumes. Accordingly, it offers advantages over previous approaches that utilize manual identification of various organs in images, and which only classify lesions using a small, finite, number of enumerated values. Accordingly, this approach provides more detailed information that can be utilized to manage treatment strategy and track disease progression and/or treatment efficacy over time in a finely grained and accurate fashion.

[0302] The LPI for each detected hotspot was computed based on reference values determined from blood pool (as measured from an identified aorta volume) and liver reference regions, e.g., similar to that described in Example 7 above. In particular, a blood pool reference intensity (SUV value) was measured using an aorta volume within the PET image, and a liver reference intensity was measured from a liver volume identified in the PET image. Both the blood pool and liver reference intensity levels were measured as the mean SUV (SUV_{mean}) within the corresponding volume in a PET image. Reference LPI index values of 1 and 2 were assigned to the blood pool and liver reference intensities, respectively. A maximal reference index level of 3 was assigned to a reference intensity corresponding to twice the liver reference intensity value.

[0303] Each individual detected hotspots was assigned LPI scores based (i) a measured hotspot intensity value for the hotspot and (ii) comparison with the blood pool and liver reference intensity values and the aforementioned reference index levels. In particular, for each hotspot, an individual hotspot intensity value corresponding to a mean lesion uptake was calculated as a mean SUV across voxels of the hotspot. For an individual hotspot, an LPI equal to 1 was assigned if the mean lesion standard uptake value (SUV_{mean}) equal to the blood pool reference value; an LPI equal to 2 was assigned for hotspots having mean lesion uptake values equal to the liver reference uptake value; and an LPI equal to 3 was assigned to hotspots having a mean lesion uptake value is equal to or above twice the liver reference uptake. For hotspots having intensities in falling in between reference intensity values, individual hotspot LPIs were interpolated from based on the hotspot intensity value and reference intensity-index value pairs.

[0304] An aggregated lesion volume within the bone region was computed as a weighted sum of volumes of individual hotspots detected within the bone volume in the PET image, with the volume of each hotspot weighted by its corresponding LPI. In this manner a PSMA-weighted Total Lesion Volume (PLTV) index was computed for the bone region, denoted $PTLV_{bone}$. PLTV index values were also computed for the lymph ($PTLV_{lymph}$) and prostate ($PTLV_{prostate}$) regions analogously. For each of the three regions, PET/CT image data sets for subjects having various different indications were used to automatically determine PLTV index values.

[0305] Performance of the AI-based automated technique was evaluated based on comparison with manual organ segmentation and annotation of images to identify hotspots, and PLTV index values computed for different indications were compared across regions.

[0306] Using manual expert interpretation as a gold standard for comparison, the automated hotspot detection algorithm was determined to have sensitivity of 92.1% for bone lesions (97.2% in the 52 automatically segmented bones) and 96.2% for lymph lesions. On average, 17 bone hotspots were detected per scan and 23 lymph hotspots were detected per scan.

[0307] PLTV index values were computed using a sample dataset of 197 PET/CT images for the bone region, 99 PET/CT image for the lymph region, and 43 PET/CT images for the prostate region. In the data sets, 94% of individual hotspot LPI's determined were between 1 and 3, with minimum LPIs determined for bone, lymph, and prostate regions of 0.82, 1.1, and 1.3, respectively. Median LPI values for the bone, lymph, and prostate regions were 1.5, 2.0, and 2.3, respectively.

[0308] For each region, PLTV indices were computed for subjects having various different indications as follows: Treatment Response (TR), Screening (S), Newly Diagnosed (ND), Metastatic (M), Suspected Recurrence (SR), Recurrent (R). PLTV index values were compared across indications by calculating the means of the values in the interquartile range (IQR_{mean}), hence excluding outliers. Ordering indications by IQR_{mean} of the PLTV values yields $TR < S < ND < R < SR < M$ for the bone region, $S = M < R < ND < SR$ for lymph region, and $M < SR < R < S < ND$ for prostate region, aligning well with clinical expectations of disease state.

[0309] FIG. 19 shows PLTV index values determined for the various indications for each region. FIG. 20 shows a scatter plot of individual hotspot LPI values for hotspots detected in each of the three regions.

[0310] In comparison with manual approaches for determining reference values for grading lesions, such as the of *Eiber et al., 2017*, the blood and liver automated reference values are based on a larger image volume compared to the manual reference method and are hence expected to be more robust. FIG. 21 compares reference values computed via the automated segmentation approaches described herein for the blood pool (left graph) and liver (right graph) regions with reference values computed using manual identification of organ boundaries. As shown in FIG. 21, in a sample of 20 PET/CT scans, the correlation (Pearson's r) between automated and manual reference was 0.92 for blood and 0.87 for liver.

[0311] Accordingly, this example demonstrates use of the automated segmentation, hotspot detection, and index value calculation for detecting cancer and tracking cancer progression and/or response to treatment over time.

D. Imaging Agents

[0312] In certain embodiments, 3D functional images are nuclear medicine images that use imaging agents comprising radiopharmaceuticals. Nuclear medicine images are obtained following administration of a radiopharmaceutical to a patient, and provide information regarding the distribution of the radiopharmaceutical within the patient. Radiopharmaceuticals are compounds that comprise a radionuclide.

[0313] Nuclear medicine images (e.g., PET scans; e.g., SPECT scans; e.g., whole-body scans; e.g. composite PET-CT images; e.g., composite SPECT-CT images) detect radiation emitted from the radionuclides of radiopharmaceuticals to form an image. The distribution of a particular radiopharmaceutical within a patient may be determined by biological mechanisms such as blood flow or perfusion, as well as by specific enzymatic or receptor binding interactions. Different radiopharmaceuticals may be designed to take advantage of different biological mechanisms and/or particular specific enzymatic or receptor binding interactions and thus, when administered to a patient, selectively concentrate within particular types of tissue and/or regions within the patient. Greater amounts of radiation are emitted from regions within the patient that have higher concentrations of radiopharmaceutical than other regions, such that these regions appear brighter in nuclear medicine images. Accordingly, intensity variations within a nuclear medicine image can be used to map the distribution of radiopharmaceutical within the patient. This mapped distribution of radiopharmaceutical within the patient can be used to, for example, infer the presence of cancerous tissue within various regions of the patient's body.

[0314] For example, upon administration to a patient, technetium 99m methylenediphosphonate (^{99m}Tc MDP) selectively accumulates within the skeletal region of the patient, in particular at sites with abnormal osteogenesis associated with malignant bone lesions. The selective concentration of radiopharmaceutical at these sites produces identifiable hotspots – localized regions of high intensity in nuclear medicine images. Accordingly, presence of malignant bone lesions associated with metastatic prostate cancer can be inferred by identifying such hotspots within a whole-body scan of the patient. Risk indices that correlate with patient overall survival and other prognostic metrics indicative of disease state, progression, treatment efficacy, and the like, can be computed based on automated analysis of intensity variations in whole-body scans obtained following administration of ^{99m}Tc MDP to a patient. In certain embodiments, other radiopharmaceuticals can also be used in a similar fashion to ^{99m}Tc MDP.

[0315] In certain embodiments, the particular radiopharmaceutical used depends on the particular nuclear medicine imaging modality used. For example ^{18}F sodium fluoride (NaF) also accumulates in bone lesions, similar to ^{99m}Tc MDP, but can be used with PET imaging. In certain embodiments, PET imaging may also utilize a radioactive form of the vitamin choline, which is readily absorbed by prostate cancer cells.

[0316] In certain embodiments, radiopharmaceuticals that selectively bind to particular proteins or receptors of interest – particularly those whose expression is increased in cancerous tissue may be used. Such proteins or receptors of interest include, but are not limited to tumor antigens, such as CEA, which is expressed in colorectal carcinomas, Her2/neu, which is expressed in

multiple cancers, BRCA 1 and BRCA 2, expressed in breast and ovarian cancers; and TRP-1 and -2, expressed in melanoma.

[0317] For example, human prostate-specific membrane antigen (PSMA) is upregulated in prostate cancer, including metastatic disease. PSMA is expressed by virtually all prostate cancers and its expression is further increased in poorly differentiated, metastatic and hormone refractory carcinomas. Accordingly, radiopharmaceuticals corresponding to PSMA binding agents (e.g., compounds that a high affinity to PSMA) labelled with one or more radionuclide(s) can be used to obtain nuclear medicine images of a patient from which the presence and/or state of prostate cancer within a variety of regions (e.g., including, but not limited to skeletal regions) of the patient can be assessed. In certain embodiments, nuclear medicine images obtained using PSMA binding agents are used to identify the presence of cancerous tissue within the prostate, when the disease is in a localized state. In certain embodiments, nuclear medicine images obtained using radiopharmaceuticals comprising PSMA binding agents are used to identify the presence of cancerous tissue within a variety of regions that include not only the prostate, but also other organs and tissue regions such as lungs, lymph nodes, and bones, as is relevant when the disease is metastatic.

[0318] In particular, upon administration to a patient, radionuclide labelled PSMA binding agents selectively accumulate within cancerous tissue, based on their affinity to PSMA. In a similar manner to that described above with regard to ^{99m}Tc MDP, the selective concentration of radionuclide labelled PSMA binding agents at particular sites within the patient produces detectable hotspots in nuclear medicine images. As PSMA binding agents concentrate within a variety of cancerous tissues and regions of the body expressing PSMA, localized cancer within a prostate of the patient and/or metastatic cancer in various regions of the patient's body can be detected, and evaluated. As described in the following, risk indices that correlate with patient overall survival and other prognostic metrics indicative of disease state, progression, treatment efficacy, and the like, can be computed based on automated analysis of intensity variations in nuclear medicine images obtained following administration of a PSMA binding agent radiopharmaceutical to a patient.

[0319] A variety of radionuclide labelled PSMA binding agents may be used as radiopharmaceutical imaging agents for nuclear medicine imaging to detect and evaluate prostate cancer. In certain embodiments, the particular radionuclide labelled PSMA binding agent that is used depends on factors such as the particular imaging modality (e.g., PET; e.g., SPECT) and the particular regions (e.g., organs) of the patient to be imaged. For example, certain radionuclide labelled PSMA binding agents are suited for PET imaging, while others are suited for SPECT imaging. For example, certain radionuclide labelled PSMA binding agents facilitate imaging a prostate of the


patient, and are used primarily when the disease is localized, while others facilitate imaging organs and regions throughout the patient's body, and are useful for evaluating metastatic prostate cancer.

[0320] A variety of PSMA binding agents and radionuclide labelled versions thereof are described in U.S. Patent Nos. 8,778,305, 8,211,401, and 8,962,799, each of which are incorporated herein by reference in their entireties.

i. PET imaging radionuclide labelled PSMA binding agents

[0321] In certain embodiments, the radionuclide labelled PSMA binding agent is a radionuclide labelled PSMA binding agent appropriate for PET imaging.


[0322] In certain embodiments, the radionuclide labelled PSMA binding agent comprises [18F]DCFPyL (also referred to as PyLTM; also referred to as DCFPyL-18F):



[18F]DCFPyL,

or a pharmaceutically acceptable salt thereof.


[0323] In certain embodiments, the radionuclide labelled PSMA binding agent comprises [18F]DCFBC:



[18F]DCFBC,


or a pharmaceutically acceptable salt thereof.

[0324] In certain embodiments, the radionuclide labelled PSMA binding agent comprises ⁶⁸Ga-PSMA-HBED-CC (also referred to as ⁶⁸Ga-PSMA-11):

⁶⁸Ga-PSMA-HBED-CC,

or a pharmaceutically acceptable salt thereof.


[0325] In certain embodiments, the radionuclide labelled PSMA binding agent comprises PSMA-617:



PSMA-617,

or a pharmaceutically acceptable salt thereof. In certain embodiments, the radionuclide labelled PSMA binding agent comprises ⁶⁸Ga-PSMA-617, which is PSMA-617 labelled with ⁶⁸Ga, or a pharmaceutically acceptable salt thereof. In certain embodiments, the radionuclide labelled PSMA binding agent comprises ¹⁷⁷Lu-PSMA-617, which is PSMA-617 labelled with ¹⁷⁷Lu, or a pharmaceutically acceptable salt thereof.


[0326] In certain embodiments, the radionuclide labelled PSMA binding agent comprises PSMA-I&T:



PSMA-I&T,

or a pharmaceutically acceptable salt thereof. In certain embodiments, the radionuclide labelled PSMA binding agent comprises ^{68}Ga -PSMA-I&T, which is PSMA-I&T labelled with ^{68}Ga , or a pharmaceutically acceptable salt thereof.

[0327] In certain embodiments, the radionuclide labelled PSMA binding agent comprises PSMA-1007:




PSMA-1007,

or a pharmaceutically acceptable salt thereof. In certain embodiments, the radionuclide labelled PSMA binding agent comprises ^{18}F -PSMA-1007, which is PSMA-1007 labelled with ^{18}F , or a pharmaceutically acceptable salt thereof.

ii. SPECT imaging radionuclide labelled PSMA binding agents

[0328] In certain embodiments, the radionuclide labelled PSMA binding agent is a radionuclide labelled PSMA binding agent appropriate for SPECT imaging.


[0329] In certain embodiments, the radionuclide labelled PSMA binding agent comprises 1404 (also referred to as MIP-1404):



1404,

or a pharmaceutically acceptable salt thereof.


[0330] In certain embodiments, the radionuclide labelled PSMA binding agent comprises 1405 (also referred to as MIP-1405):



1405,

or a pharmaceutically acceptable salt thereof.


[0331] In certain embodiments, the radionuclide labelled PSMA binding agent comprises 1427 (also referred to as MIP-1427):



1427,

or a pharmaceutically acceptable salt thereof.

[0332] In certain embodiments, the radionuclide labelled PSMA binding agent comprises 1428 (also referred to as MIP-1428):




1428,

or a pharmaceutically acceptable salt thereof.

[0333] In certain embodiments, the PSMA binding agent is labelled with a radionuclide by chelating it to a radioisotope of a metal [e.g., a radioisotope of technetium (Tc) (e.g., technetium-99m (^{99m}Tc)); e.g., a radioisotope of rhenium (Re) (e.g., rhenium-188 (^{188}Re); e.g., rhenium-186 (^{186}Re)); e.g., a radioisotope of yttrium (Y) (e.g., ^{90}Y); e.g., a radioisotope of lutetium (Lu) (e.g., ^{177}Lu); e.g., a radioisotope of gallium (Ga) (e.g., ^{68}Ga ; e.g., ^{67}Ga); e.g., a radioisotope of indium (e.g., ^{111}In); e.g., a radioisotope of copper (Cu) (e.g., ^{67}Cu)].


[0334] In certain embodiments, 1404 is labelled with a radionuclide (e.g., chelated to a radioisotope of a metal). In certain embodiments, the radionuclide labelled PSMA binding agent comprises ^{99m}Tc -MIP-1404, which is 1404 labelled with (e.g., chelated to) ^{99m}Tc :



^{99m}Tc -MIP-1404,


or a pharmaceutically acceptable salt thereof. In certain embodiments, 1404 may be chelated to other metal radioisotopes [e.g., a radioisotope of rhenium (Re) (e.g., rhenium-188 (^{188}Re); e.g., rhenium-186 (^{186}Re)); e.g., a radioisotope of yttrium (Y) (e.g., ^{90}Y); e.g., a radioisotope of lutetium (Lu) (e.g., ^{177}Lu); e.g., a radioisotope of gallium (Ga) (e.g., ^{68}Ga ; e.g., ^{67}Ga); e.g., a radioisotope of indium (e.g., ^{111}In); e.g., a radioisotope of copper (Cu) (e.g., ^{67}Cu)] to form a compound having a structure similar to the structure shown above for ^{99m}Tc -MIP-1404, with the other metal radioisotope substituted for ^{99m}Tc .

[0335] In certain embodiments, 1405 is labelled with a radionuclide (e.g., chelated to a radioisotope of a metal). In certain embodiments, the radionuclide labelled PSMA binding agent comprises ^{99m}Tc -MIP-1405, which is 1405 labelled with (e.g., chelated to) ^{99m}Tc :

^{99m}Tc-MIP-1405,

or a pharmaceutically acceptable salt thereof. In certain embodiments, 1405 may be chelated to other metal radioisotopes [e.g., a radioisotope of rhenium (Re) (e.g., rhenium-188 (¹⁸⁸Re); e.g., rhenium-186 (¹⁸⁶Re)); e.g., a radioisotope of yttrium (Y) (e.g., ⁹⁰Y); e.g., a radioisotope of lutetium (Lu) (e.g., ¹⁷⁷Lu); e.g., a radioisotope of gallium (Ga) (e.g., ⁶⁸Ga; e.g., ⁶⁷Ga); e.g., a radioisotope of indium (e.g., ¹¹¹In); e.g., a radioisotope of copper (Cu) (e.g., ⁶⁷Cu)] to form a compound having a structure similar to the structure shown above for ^{99m}Tc-MIP-1405, with the other metal radioisotope substituted for ^{99m}Tc.


[0336] In certain embodiments, 1427 is labelled with (e.g., chelated to) a radioisotope of a metal, to form a compound according to the formula below:



1427 chelated to a metal,

or a pharmaceutically acceptable salt thereof, wherein M is a metal radioisotope [e.g., a radioisotope of technetium (Tc) (e.g., technetium-99m (^{99m}Tc)); e.g., a radioisotope of rhenium (Re) (e.g., rhenium-188 (¹⁸⁸Re); e.g., rhenium-186 (¹⁸⁶Re)); e.g., a radioisotope of yttrium (Y) (e.g., ⁹⁰Y); e.g., a radioisotope of lutetium (Lu) (e.g., ¹⁷⁷Lu); e.g., a radioisotope of gallium (Ga) (e.g., ⁶⁸Ga; e.g., ⁶⁷Ga); e.g., a radioisotope of indium (e.g., ¹¹¹In); e.g., a radioisotope of copper (Cu) (e.g., ⁶⁷Cu)] with which 1427 is labelled.


[0337] In certain embodiments, 1428 is labelled with (e.g., chelated to) a radioisotope of a metal, to form a compound according to the formula below:



1428 chelated to a metal,

or a pharmaceutically acceptable salt thereof, wherein M is a metal radioisotope [e.g., a radioisotope of technetium (Tc) (e.g., technetium-99m (^{99m}Tc)); e.g., a radioisotope of rhenium (Re) (e.g., rhenium-188 (^{188}Re); e.g., rhenium-186 (^{186}Re)); e.g., a radioisotope of yttrium (Y) (e.g., ^{90}Y); e.g., a radioisotope of lutetium (Lu) (e.g., ^{177}Lu); e.g., a radioisotope of gallium (Ga) (e.g., ^{68}Ga ; e.g., ^{67}Ga); e.g., a radioisotope of indium (e.g., ^{111}In); e.g., a radioisotope of copper (Cu) (e.g., ^{67}Cu)] with which 1428 is labelled.

[0338] In certain embodiments, the radionuclide labelled PSMA binding agent comprises PSMA I&S:



PSMA I&S,

or a pharmaceutically acceptable salt thereof. In certain embodiments, the radionuclide labelled PSMA binding agent comprises ^{99m}Tc -PSMA I&S, which is PSMA I&S labelled with ^{99m}Tc , or a pharmaceutically acceptable salt thereof.

E. Computer System and Network Architecture

[0339] As shown in FIG. 17, an implementation of a network environment 1700 for use in providing systems, methods, and architectures described herein is shown and described. In brief

overview, referring now to FIG. 17, a block diagram of an exemplary cloud computing environment 1700 is shown and described. The cloud computing environment 1700 may include one or more resource providers 1702a, 1702b, 1702c (collectively, 1702). Each resource provider 1702 may include computing resources. In some implementations, computing resources may include any hardware and/or software used to process data. For example, computing resources may include hardware and/or software capable of executing algorithms, computer programs, and/or computer applications. In some implementations, exemplary computing resources may include application servers and/or databases with storage and retrieval capabilities. Each resource provider 1702 may be connected to any other resource provider 1702 in the cloud computing environment 1700. In some implementations, the resource providers 1702 may be connected over a computer network 1708. Each resource provider 1702 may be connected to one or more computing device 1704a, 1704b, 1704c (collectively, 1704), over the computer network 1708.

[0340] The cloud computing environment 1700 may include a resource manager 1706. The resource manager 1706 may be connected to the resource providers 1702 and the computing devices 1704 over the computer network 1708. In some implementations, the resource manager 1706 may facilitate the provision of computing resources by one or more resource providers 1702 to one or more computing devices 1704. The resource manager 1706 may receive a request for a computing resource from a particular computing device 1704. The resource manager 1706 may identify one or more resource providers 1702 capable of providing the computing resource requested by the computing device 1704. The resource manager 1706 may select a resource provider 1702 to provide the computing resource. The resource manager 1706 may facilitate a connection between the resource provider 1702 and a particular computing device 1704. In some implementations, the resource manager 1706 may establish a connection between a particular resource provider 1702 and a particular computing device 1704. In some implementations, the resource manager 1706 may redirect a particular computing device 1704 to a particular resource provider 1702 with the requested computing resource.

[0341] FIG. 18 shows an example of a computing device 1800 and a mobile computing device 1850 that can be used to implement the techniques described in this disclosure. The computing device 1800 is intended to represent various forms of digital computers, such as laptops, desktops, workstations, personal digital assistants, servers, blade servers, mainframes, and other appropriate computers. The mobile computing device 1850 is intended to represent various forms of mobile devices, such as personal digital assistants, cellular telephones, smart-phones, and other similar computing devices. The components shown here, their connections and relationships, and their functions, are meant to be examples only, and are not meant to be limiting.

[0342] The computing device 1800 includes a processor 1802, a memory 1804, a storage device 1806, a high-speed interface 1808 connecting to the memory 1804 and multiple high-speed expansion ports 1810, and a low-speed interface 1812 connecting to a low-speed expansion port 1814 and the storage device 1806. Each of the processor 1802, the memory 1804, the storage device 1806, the high-speed interface 1808, the high-speed expansion ports 1810, and the low-speed interface 1812, are interconnected using various busses, and may be mounted on a common motherboard or in other manners as appropriate. The processor 1802 can process instructions for execution within the computing device 1800, including instructions stored in the memory 1804 or on the storage device 1806 to display graphical information for a GUI on an external input/output device, such as a display 1816 coupled to the high-speed interface 1808. In other implementations, multiple processors and/or multiple buses may be used, as appropriate, along with multiple memories and types of memory. Also, multiple computing devices may be connected, with each device providing portions of the necessary operations (e.g., as a server bank, a group of blade servers, or a multi-processor system). Thus, as the term is used herein, where a plurality of functions are described as being performed by “a processor”, this encompasses embodiments wherein the plurality of functions are performed by any number of processors (one or more) of any number of computing devices (one or more). Furthermore, where a function is described as being performed by “a processor”, this encompasses embodiments wherein the function is performed by any number of processors (one or more) of any number of computing devices (one or more) (e.g., in a distributed computing system).

[0343] The memory 1804 stores information within the computing device 1800. In some implementations, the memory 1804 is a volatile memory unit or units. In some implementations, the memory 1804 is a non-volatile memory unit or units. The memory 1804 may also be another form of computer-readable medium, such as a magnetic or optical disk.

[0344] The storage device 1806 is capable of providing mass storage for the computing device 1800. In some implementations, the storage device 1806 may be or contain a computer-readable medium, such as a floppy disk device, a hard disk device, an optical disk device, or a tape device, a flash memory or other similar solid state memory device, or an array of devices, including devices in a storage area network or other configurations. Instructions can be stored in an information carrier. The instructions, when executed by one or more processing devices (for example, processor 1802), perform one or more methods, such as those described above. The instructions can also be stored by one or more storage devices such as computer- or machine-readable mediums (for example, the memory 1804, the storage device 1806, or memory on the processor 1802).

[0345] The high-speed interface 1808 manages bandwidth-intensive operations for the computing device 1800, while the low-speed interface 1812 manages lower bandwidth-intensive operations. Such allocation of functions is an example only. In some implementations, the high-speed interface 1808 is coupled to the memory 1804, the display 1816 (e.g., through a graphics processor or accelerator), and to the high-speed expansion ports 1810, which may accept various expansion cards (not shown). In the implementation, the low-speed interface 1812 is coupled to the storage device 1806 and the low-speed expansion port 1814. The low-speed expansion port 1814, which may include various communication ports (e.g., USB, Bluetooth®, Ethernet, wireless Ethernet) may be coupled to one or more input/output devices, such as a keyboard, a pointing device, a scanner, or a networking device such as a switch or router, e.g., through a network adapter.

[0346] The computing device 1800 may be implemented in a number of different forms, as shown in the figure. For example, it may be implemented as a standard server 1820, or multiple times in a group of such servers. In addition, it may be implemented in a personal computer such as a laptop computer 1822. It may also be implemented as part of a rack server system 1824. Alternatively, components from the computing device 1800 may be combined with other components in a mobile device (not shown), such as a mobile computing device 1850. Each of such devices may contain one or more of the computing device 1800 and the mobile computing device 1850, and an entire system may be made up of multiple computing devices communicating with each other.

[0347] The mobile computing device 1850 includes a processor 1852, a memory 1864, an input/output device such as a display 1854, a communication interface 1866, and a transceiver 1868, among other components. The mobile computing device 1850 may also be provided with a storage device, such as a micro-drive or other device, to provide additional storage. Each of the processor 1852, the memory 1864, the display 1854, the communication interface 1866, and the transceiver 1868, are interconnected using various buses, and several of the components may be mounted on a common motherboard or in other manners as appropriate.

[0348] The processor 1852 can execute instructions within the mobile computing device 1850, including instructions stored in the memory 1864. The processor 1852 may be implemented as a chipset of chips that include separate and multiple analog and digital processors. The processor 1852 may provide, for example, for coordination of the other components of the mobile computing device 1850, such as control of user interfaces, applications run by the mobile computing device 1850, and wireless communication by the mobile computing device 1850.

[0349] The processor 1852 may communicate with a user through a control interface 1858 and a display interface 1856 coupled to the display 1854. The display 1854 may be, for example, a

TFT (Thin-Film-Transistor Liquid Crystal Display) display or an OLED (Organic Light Emitting Diode) display, or other appropriate display technology. The display interface 1856 may comprise appropriate circuitry for driving the display 1854 to present graphical and other information to a user. The control interface 1858 may receive commands from a user and convert them for submission to the processor 1852. In addition, an external interface 1862 may provide communication with the processor 1852, so as to enable near area communication of the mobile computing device 1850 with other devices. The external interface 1862 may provide, for example, for wired communication in some implementations, or for wireless communication in other implementations, and multiple interfaces may also be used.

[0350] The memory 1864 stores information within the mobile computing device 1850. The memory 1864 can be implemented as one or more of a computer-readable medium or media, a volatile memory unit or units, or a non-volatile memory unit or units. An expansion memory 1874 may also be provided and connected to the mobile computing device 1850 through an expansion interface 1872, which may include, for example, a SIMM (Single In Line Memory Module) card interface. The expansion memory 1874 may provide extra storage space for the mobile computing device 1850, or may also store applications or other information for the mobile computing device 1850. Specifically, the expansion memory 1874 may include instructions to carry out or supplement the processes described above, and may include secure information also. Thus, for example, the expansion memory 1874 may be provide as a security module for the mobile computing device 1850, and may be programmed with instructions that permit secure use of the mobile computing device 1850. In addition, secure applications may be provided via the SIMM cards, along with additional information, such as placing identifying information on the SIMM card in a non-hackable manner.

[0351] The memory may include, for example, flash memory and/or NVRAM memory (non-volatile random access memory), as discussed below. In some implementations, instructions are stored in an information carrier. that the instructions, when executed by one or more processing devices (for example, processor 1852), perform one or more methods, such as those described above. The instructions can also be stored by one or more storage devices, such as one or more computer- or machine-readable mediums (for example, the memory 1864, the expansion memory 1874, or memory on the processor 1852). In some implementations, the instructions can be received in a propagated signal, for example, over the transceiver 1868 or the external interface 1862.

[0352] The mobile computing device 1850 may communicate wirelessly through the communication interface 1866, which may include digital signal processing circuitry where

necessary. The communication interface 1866 may provide for communications under various modes or protocols, such as GSM voice calls (Global System for Mobile communications), SMS (Short Message Service), EMS (Enhanced Messaging Service), or MMS messaging (Multimedia Messaging Service), CDMA (code division multiple access), TDMA (time division multiple access), PDC (Personal Digital Cellular), WCDMA (Wideband Code Division Multiple Access), CDMA2000, or GPRS (General Packet Radio Service), among others. Such communication may occur, for example, through the transceiver 1868 using a radio-frequency. In addition, short-range communication may occur, such as using a Bluetooth®, Wi-Fi™, or other such transceiver (not shown). In addition, a GPS (Global Positioning System) receiver module 1870 may provide additional navigation- and location-related wireless data to the mobile computing device 1850, which may be used as appropriate by applications running on the mobile computing device 1850.

[0353] The mobile computing device 1850 may also communicate audibly using an audio codec 1860, which may receive spoken information from a user and convert it to usable digital information. The audio codec 1860 may likewise generate audible sound for a user, such as through a speaker, e.g., in a handset of the mobile computing device 1850. Such sound may include sound from voice telephone calls, may include recorded sound (e.g., voice messages, music files, etc.) and may also include sound generated by applications operating on the mobile computing device 1850.

[0354] The mobile computing device 1850 may be implemented in a number of different forms, as shown in the figure. For example, it may be implemented as a cellular telephone 1880. It may also be implemented as part of a smart-phone 1882, personal digital assistant, or other similar mobile device.

[0355] Various implementations of the systems and techniques described here can be realized in digital electronic circuitry, integrated circuitry, specially designed ASICs (application specific integrated circuits), computer hardware, firmware, software, and/or combinations thereof. These various implementations can include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which may be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device.

[0356] These computer programs (also known as programs, software, software applications or code) include machine instructions for a programmable processor, and can be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the terms machine-readable medium and computer-readable medium refer to any computer program product, apparatus and/or device (e.g., magnetic discs, optical disks,

memory, Programmable Logic Devices (PLDs)) used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine instructions as a machine-readable signal. The term machine-readable signal refers to any signal used to provide machine instructions and/or data to a programmable processor.

[0357] To provide for interaction with a user, the systems and techniques described here can be implemented on a computer having a display device (e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor) for displaying information to the user and a keyboard and a pointing device (e.g., a mouse or a trackball) by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback (e.g., visual feedback, auditory feedback, or tactile feedback); and input from the user can be received in any form, including acoustic, speech, or tactile input.

[0358] The systems and techniques described here can be implemented in a computing system that includes a back end component (e.g., as a data server), or that includes a middleware component (e.g., an application server), or that includes a front end component (e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the systems and techniques described here), or any combination of such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication (e.g., a communication network). Examples of communication networks include a local area network (LAN), a wide area network (WAN), and the Internet.

[0359] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0360] In some implementations, the various modules described herein can be separated, combined or incorporated into single or combined modules. The modules depicted in the figures are not intended to limit the systems described herein to the software architectures shown therein.

[0361] Elements of different implementations described herein may be combined to form other implementations not specifically set forth above. Elements may be left out of the processes, computer programs, databases, etc. described herein without adversely affecting their operation. In addition, the logic flows depicted in the figures do not require the particular order shown, or sequential order, to achieve desirable results. Various separate elements may be combined into one or more individual elements to perform the functions described herein.

[0362] Throughout the description, where apparatus and systems are described as having, including, or comprising specific components, or where processes and methods are described as having, including, or comprising specific steps, it is contemplated that, additionally, there are apparatus, and systems of the present invention that consist essentially of, or consist of, the recited components, and that there are processes and methods according to the present invention that consist essentially of, or consist of, the recited processing steps.

[0363] It should be understood that the order of steps or order for performing certain action is immaterial so long as the invention remains operable. Moreover, two or more steps or actions may be conducted simultaneously.

[0364] While the invention has been particularly shown and described with reference to specific preferred embodiments, it should be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention as defined by the appended claims.

[0365] While the invention has been particularly shown and described with reference to specific preferred embodiments, it should be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention as defined by the appended claims.

[0366] Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

[0367] The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any form of suggestion that the prior art forms part of the common general knowledge in Australia.

CLAIMS

1. A method for automatically processing 3D images to automatically identify cancerous lesions within a subject, the method comprising:

(a) receiving, by a processor of a computing device, a 3D anatomical image of the subject obtained using an anatomical imaging modality, wherein the 3D anatomical image comprises a graphical representation of tissue within the subject;

(b) automatically identifying, by the processor, using one or more machine learning modules, for each of a plurality of target tissue regions, a corresponding target volume of interest (VOI) within the 3D anatomical image;

(c) determining, by the processor, a 3D segmentation map representing a plurality of 3D segmentation masks, each 3D segmentation mask representing a particular identified target VOI;

(d) receiving, by the processor, a 3D functional image of the subject obtained using a functional imaging modality;

(e) identifying, within the 3D functional image, one or more 3D volume(s), each corresponding to an identified target VOI, using the 3D segmentation map; and

(f) automatically detecting, by the processor, within at least a portion of the one or more 3D volumes identified within the 3D functional image, one or more hotspots determined to represent lesions based on intensities of voxels within the 3D functional image.

2. The method of claim 1, comprising using, by the processor, the one or more detected hotspots to determine a cancer status for the subject.

3. The method of either of claims 1 or 2, wherein the target tissue regions comprise one or more reference tissue regions and wherein the method comprises:

using the 3D segmentation map to identify, by the processor, within the 3D functional image, one or more 3D reference volume(s), each corresponding to a particular reference tissue region;

determining, by the processor, one or more reference intensity values, each associated with a particular 3D reference volume of the one or more 3D reference volume(s) and corresponding to a measure of intensity within the particular 3D reference volume;

determining, by the processor, one or more individual hotspot intensity values, each associated with a particular hotspot of at least a portion of the detected one or more hotspots and corresponding to a measure of intensity of the particular hotspot; and

determining, by the processor, one or more individual hotspot index values using the one or more individual hotspot intensity values and the one or more reference intensity values.

4. The method of claim 3, wherein the reference tissue regions comprise one or more members selected from the group consisting of: a liver, an aorta, and a parotid gland.

5. The method of any one of claims 2 to 4, comprising, determining, by the processor, an overall index value indicative of a cancer status of the subject using at least a portion of the one or more hotspot index values.

6. The method of claim 5, wherein the overall index value is determined as a weighted sum of at least a portion of the individual hotspot index values.

7. The method of any one of claims 1 to 6, wherein:
the 3D anatomical image is an *x*-ray computed tomography (CT) image, and
the 3D functional image is a 3D positron emission tomography (PET) image.
8. The method of claim 7, wherein the 3D PET image of the subject is obtained following administration to the subject of a radiopharmaceutical comprising a prostate-specific membrane antigen (PSMA) binding agent.
9. The method of claim 8, wherein the radiopharmaceutical comprises [¹⁸F]DCFPyL.
10. A system for automatically processing 3D images to automatically identify cancerous lesions within a subject, the system comprising:
a processor of a computing device; and
a memory having instructions stored thereon, wherein the instructions, when executed by the processor, cause the processor to:
 - (a) receive a 3D anatomical image of the subject obtained using an anatomical imaging modality, wherein the 3D anatomical image comprises a graphical representation of tissue within the subject;
 - (b) automatically identify, using one or more machine learning modules, for each of a plurality of target tissue regions, a corresponding target volume of interest (VOI) within the 3D anatomical image;
 - (c) determine a 3D segmentation map representing a plurality of 3D segmentation masks, each 3D segmentation mask representing a particular identified target VOI;

- (d) receive a 3D functional image of the subject obtained using a functional imaging modality;
- (e) identify, within the 3D functional image, one or more 3D volume(s), each corresponding to an identified target VOI, using the 3D segmentation map; and
- (f) automatically detect, within at least a portion of the one or more 3D volumes identified within the 3D functional image, one or more hotspots determined to represent lesions based on intensities of voxels within the 3D functional image.

11. The system of claim 10, wherein the instructions cause the processor to use the one or more detected hotspots to determine a cancer status for the subject.

12. The system of either of claims 10 or 11, wherein the target tissue regions comprise one or more reference tissue regions and wherein the instructions cause the processor to:

use the 3D segmentation map to identify, within the 3D functional image, one or more 3D reference volume(s), each corresponding to a particular reference tissue region;

determine one or more reference intensity values, each associated with a particular 3D reference volume of the one or more 3D reference volume(s) and corresponding to a measure of intensity within the particular 3D reference volume;

determine one or more individual hotspot intensity values, each associated with a particular hotspot of at least a portion of the detected one or more hotspots and corresponding to a measure of intensity of the particular hotspot; and

determine one or more individual hotspot index values using the one or more individual hotspot intensity values and the one or more reference intensity values.

13. The system of claim 12, wherein the reference tissue regions comprise one or more members selected from the group consisting of: a liver, an aorta, and a parotid gland.
14. The system of any one of claims 11 to 13, wherein the instructions cause the processor to determine an overall index value indicative of a cancer status of the subject using at least a portion of the one or more hotspot index values.
15. The system of claim 14, wherein the overall index value is determined as a weighted sum of at least a portion of the individual hotspot index values.
16. The system of any one of claims 10 to 15, wherein:
the 3D anatomical image is an *x*-ray computed tomography (CT) image, and
the 3D functional image is a 3D positron emission tomography (PET) image.
17. The system of claim 16, wherein the 3D PET image of the subject is obtained following administration to the subject of a radiopharmaceutical comprising a prostate-specific membrane antigen (PSMA) binding agent.
18. The system of claim 17, wherein the radiopharmaceutical comprises [¹⁸F]DCFPyL.

19. A method for automatically processing 3D images to identify, and measure uptake of radiopharmaceutical in, cancerous lesions within a subject having or at risk for a cancer, the method comprising:

(a) receiving, by a processor of a computing device, a 3D anatomical image of the subject obtained using an anatomical imaging modality, wherein the 3D anatomical image comprises a graphical representation of tissue within the subject;

(b) automatically identifying, by the processor, using one or more machine learning modules, within the 3D anatomical image:

a first skeletal volume comprising a graphical representation of one or more bones of the subject;

a first aorta volume comprising a graphical representation of at least a portion of an aorta of the subject; and

a first liver volume comprising a graphical representation of a liver of the subject;

(c) determining, by the processor, a 3D segmentation map representing a plurality of 3D segmentation masks, including a skeletal mask representing the identified first skeletal volume, an aorta mask representing the identified first aorta volume, and a liver mask representing the identified first liver volume;

(d) receiving, by the processor, a 3D functional image of the subject obtained using a functional imaging modality;

(e) automatically identifying, within the 3D functional image, using the 3D segmentation map:

a second skeletal volume corresponding to the first identified skeletal volume, within the 3D anatomical image;

a second aorta volume corresponding to the first aorta volume, identified within the 3D anatomical image; and

a second liver volume corresponding to the first liver volume, identified within the 3D anatomical image;

(f) automatically detecting, by the processor, within the second skeletal volume, one or more hotspots determined to represent lesions based on intensities of voxels within the second skeletal volume; and

(g) determining, by the processor, for each of the one or more detected hotspots, an individual hotspot index value by:

determining an aorta reference intensity level based on a measure of intensity of voxels within the second aorta volume;

determining a liver reference intensity level based on a measure of intensity of voxels within the second liver volume; and

for each individual detected hotspot:

determining a corresponding individual hotspot intensity level based on a measure of intensity of voxels of the detected hotspot; and

determining a corresponding individual hotspot index level from the individual hotspot intensity level, the aorta reference intensity level, and the liver reference intensity level.

20. The method of claim 19, comprising determining, by the processor, an overall index value indicative of a cancer status of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots.

21. The method of either of claims 19 or 20, wherein the subject has or is at risk for prostate cancer.
22. The method of any one of claims 19 to 21, wherein:
step (b) comprises automatically identifying, within the 3D anatomical image, a first prostate volume comprising a graphical representation of a prostate of the subject,
the 3D segmentation map determined at step (c) further includes a prostate mask representing the identified first prostate volume,
step (e) comprises automatically identifying, within the 3D functional image, a second prostate volume corresponding to the first identified prostate volume, within the 3D anatomical image,
step (f) comprises automatically detecting one or more hotspots in the second prostate volume, and the method further comprises:
determining, by the processor, (i) an overall bone index value indicative of a lesion content in the one or more bones of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second skeletal volume and (ii) an overall prostate index value indicative of a lesion content in the prostate of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second prostate volume.
23. The method of any one of claims 19 to 22, wherein the subject has or is at risk for metastatic cancer.

24. A system for automatically processing 3D images to identify, and measure uptake of radiopharmaceutical in, cancerous lesions within a subject having or at risk for a cancer, the system comprising:

a processor of a computing device; and

a memory having instructions stored thereon, wherein the instructions, when executed by the processor, cause the processor to:

(a) receive a 3D anatomical image of the subject obtained using an anatomical imaging modality, wherein the 3D anatomical image comprises a graphical representation of tissue within the subject;

(b) automatically identify using one or more machine learning modules, within the 3D anatomical image:

a first skeletal volume comprising a graphical representation of one or more bones of the subject;

a first aorta volume comprising a graphical representation of at least a portion of an aorta of the subject; and

a first liver volume comprising a graphical representation of a liver of the subject;

(c) determine a 3D segmentation map representing a plurality of 3D segmentation masks, including a skeletal mask representing the identified first skeletal volume, an aorta mask representing the identified first aorta volume, and a liver mask representing the identified first liver volume;

(d) receive a 3D functional image of the subject obtained using a functional imaging modality;

(e) automatically identify, within the 3D functional image, using the 3D segmentation map:

a second skeletal volume corresponding to the first identified skeletal volume, within the 3D anatomical image;

a second aorta volume corresponding to the first aorta volume, identified within the 3D anatomical image; and

a second liver volume corresponding to the first liver volume, identified within the 3D anatomical image;

(f) automatically detect, within the second skeletal volume, one or more hotspots determined to represent lesions based on intensities of voxels within the second skeletal volume; and

(g) determine, for each of the one or more detected hotspots, an individual hotspot index value by:

determining an aorta reference intensity level based on a measure of intensity of voxels within the second aorta volume;

determining a liver reference intensity level based on a measure of intensity of voxels within the second liver volume; and

for each individual detected hotspot:

determining a corresponding individual hotspot intensity level based on a measure of intensity of voxels of the detected hotspot; and

determining a corresponding individual hotspot index level from the individual hotspot intensity level, the aorta reference intensity level, and the liver reference intensity level.

25. The system of claim 24, wherein the instructions cause the processor to determine an overall index value indicative of a cancer status of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots.

26. The system of either of claims 24 or 25, wherein the subject has or is at risk for prostate cancer.

27. The system of any one of claims 24 to 26, wherein the instructions cause the processor to:

at step (b), automatically identify, within the 3D anatomical image, a first prostate volume comprising a graphical representation of a prostate of the subject;

at step (c), include a prostate mask representing the identified first prostate volume in the determined 3D segmentation map;

at step (e), automatically identify, within the 3D functional image, a second prostate volume corresponding to the first identified prostate volume, within the 3D anatomical image,

at step (f), automatically detect one or more hotspots in the second prostate volume; and

determine (i) an overall bone index value indicative of a lesion content in the one or more bones of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second skeletal volume and (ii) an overall prostate index value indicative of a lesion content in the prostate of the subject based on the individual hotspot index values of at least a portion of the one or more detected hotspots located in the second prostate volume.

28. The system of any one of claims 24 to 27, wherein the subject has or is at risk for metastatic cancer.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5A




FIG. 5B




FIG. 6A






FIG. 6B






FIG. 9A




FIG. 9B




FIG. 9C




FIG. 9D





FIG. 10A






FIG. 11B




FIG. 11A




FIG. 12




FIG. 13A




FIG. 13B




FIG. 13C




FIG. 13D




FIG. 14B




FIG. 14A




FIG. 15B




FIG. 15A




FIG. 16B




FIG. 16A





FIG. 17



18. EG.




FIG. 19

26/26




FIG. 20




FIG. 21