



To Promote the Progress



of Science and Useful Arts

The Director

of the United States Patent and Trademark Office has received an application for a patent for a new and useful invention. The title and description of the invention are enclosed. The requirements of law have been complied with, and it has been determined that a patent on the invention shall be granted under the law.

Therefore, this United States

Patent

grants to the person(s) having title to this patent the right to exclude others from making, using, offering for sale, or selling the invention throughout the United States of America or importing the invention into the United States of America, and if the invention is a process, of the right to exclude others from using, offering for sale or selling throughout the United States of America, products made by that process, for the term set forth in 35 U.S.C. 154(a)(2) or (c)(1), subject to the payment of maintenance fees as provided by 35 U.S.C. 41(b). See the Maintenance Fee Notice on the inside of the cover.



Katherine Kelly Vidal



DIRECTOR OF THE UNITED STATES PATENT AND TRADEMARK OFFICE

Maintenance Fee Notice

If the application for this patent was filed on or after December 12, 1980, maintenance fees are due three years and six months, seven years and six months, and eleven years and six months after the date of this grant, or within a grace period of six months thereafter upon payment of a surcharge as provided by law. The amount, number and timing of the maintenance fees required may be changed by law or regulation. Unless payment of the applicable maintenance fee is received in the United States Patent and Trademark Office on or before the date the fee is due or within a grace period of six months thereafter, the patent will expire as of the end of such grace period.

Patent Term Notice

If the application for this patent was filed on or after June 8, 1995, the term of this patent begins on the date on which this patent issues and ends twenty years from the filing date of the application or, if the application contains a specific reference to an earlier filed application or applications under 35 U.S.C. 120, 121, 365(c), or 386(c), twenty years from the filing date of the earliest such application (“the twenty-year term”), subject to the payment of maintenance fees as provided by 35 U.S.C. 41(b), and any extension as provided by 35 U.S.C. 154(b) or 156 or any disclaimer under 35 U.S.C. 253.

If this application was filed prior to June 8, 1995, the term of this patent begins on the date on which this patent issues and ends on the later of seventeen years from the date of the grant of this patent or the twenty-year term set forth above for patents resulting from applications filed on or after June 8, 1995, subject to the payment of maintenance fees as provided by 35 U.S.C. 41(b) and any extension as provided by 35 U.S.C. 156 or any disclaimer under 35 U.S.C. 253.



US011734629B2

(12) **United States Patent**
Lyman et al.

(10) **Patent No.:** **US 11,734,629 B2**

(45) **Date of Patent:** ***Aug. 22, 2023**

(54) **MEDICAL SCAN LABELING QUALITY ASSURANCE SYSTEM AND METHODS FOR USE THEREWITH**

(71) Applicant: **Enlitic, Inc.**, San Francisco, CA (US)

(72) Inventors: **Kevin Lyman**, Fords, NJ (US); **Anthony Upton**, Malvern (AU); **Lionel Lints**, Oakland, CA (US); **Ben Covington**, Berkeley, CA (US); **Alexander Rhodes**, Melbourne (AU)

(73) Assignee: **Enlitic, Inc.**, Fort Collins, CO (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 85 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **17/455,066**

(22) Filed: **Nov. 16, 2021**

(65) **Prior Publication Data**

US 2022/0076793 A1 Mar. 10, 2022

Related U.S. Application Data

(63) Continuation of application No. 16/363,019, filed on Mar. 25, 2019, now Pat. No. 11,211,153.
(Continued)

(51) **Int. Cl.**

G06Q 10/0631 (2023.01)

G16H 10/60 (2018.01)

(Continued)

(52) **U.S. Cl.**

CPC **G06Q 10/06315** (2013.01); **A61B 5/7264** (2013.01); **G06F 3/0482** (2013.01);

(Continued)

(58) **Field of Classification Search**

CPC A61B 5/7264; A61B 5/055; A61B 6/032; A61B 8/4416; G06T 7/0012; G06T 7/70;

(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,466,687 B1 10/2002 Uppaluri
6,524,246 B1 2/2003 Kelly

(Continued)

FOREIGN PATENT DOCUMENTS

AU 2004316027 B2 * 7/2009 G06F 19/345
CN 106295139 A 1/2017

(Continued)

OTHER PUBLICATIONS

Yamshita, Convolutional neural networks: an overview and application in radiology, Jun. 22, 2018, Insights into Imaging, pp. 611-629. (Year: 2018).*

(Continued)

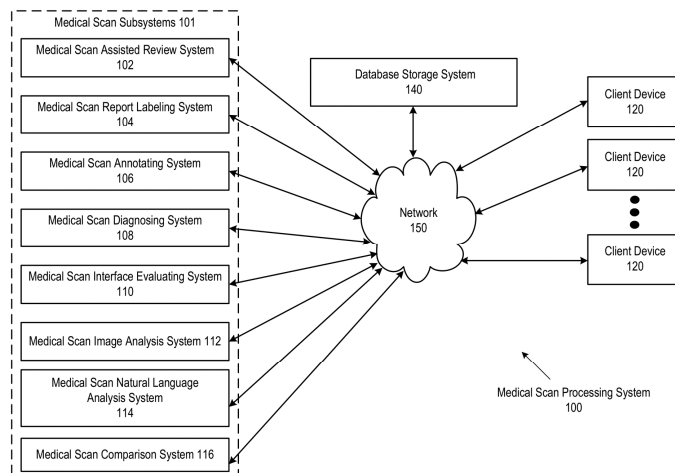
Primary Examiner — Christopher L Gilligan

(74) *Attorney, Agent, or Firm* — Cochran Freund & Young LLC

(57) **ABSTRACT**

A medical scan system is operable to receive a set of labeling data corresponding to a set of medical scans from each of a set of client devices corresponding to a set of users. The set of medical scans and each set of labeling data is transmitted to an expert client device associated with an expert user, and a set of golden labeling data and a plurality of sets of correction data are received from the expert client device. A set of performance score data is generated based on the plurality of sets of correction data, and each performance score data of the set of performance score data is assigned to a corresponding one of the set of users. An updated training set that includes the set of golden labeling data is generated, and a medical scan analysis function is retrained based on the updated training set.

20 Claims, 47 Drawing Sheets



Related U.S. Application Data

(60) Provisional application No. 62/770,334, filed on Nov. 21, 2018.

(51) **Int. Cl.**

G16H 30/40 (2018.01)
G16H 15/00 (2018.01)
G06T 5/00 (2006.01)
G06T 5/50 (2006.01)
G06T 7/00 (2017.01)
G06T 11/00 (2006.01)
G06N 5/04 (2023.01)
G16H 30/20 (2018.01)
G06N 20/00 (2019.01)
G06T 7/187 (2017.01)
G06T 7/11 (2017.01)
G06T 3/40 (2006.01)
G16H 50/20 (2018.01)
G06F 9/54 (2006.01)
G06F 3/0484 (2013.01)
A61B 5/00 (2006.01)
G06F 21/62 (2013.01)
G06Q 20/14 (2012.01)
G16H 40/20 (2018.01)
G06F 3/0484 (2022.01)
G16H 10/20 (2018.01)
G06N 5/045 (2023.01)
G06T 7/10 (2017.01)
G06T 11/20 (2006.01)
G06F 16/245 (2019.01)
G06T 7/44 (2017.01)
G06N 20/20 (2019.01)
H04L 67/12 (2022.01)
H04L 67/01 (2022.01)
G06V 10/82 (2022.01)
G06F 18/40 (2023.01)
G06F 18/214 (2023.01)
G06F 18/21 (2023.01)
G06F 18/2115 (2023.01)
G06F 18/2415 (2023.01)
G06V 10/25 (2022.01)
G06V 30/19 (2022.01)
G06V 10/764 (2022.01)
G06V 40/16 (2022.01)
G06V 10/22 (2022.01)
G16H 50/70 (2018.01)
G06T 7/70 (2017.01)
G16H 50/30 (2018.01)
A61B 5/055 (2006.01)
A61B 6/03 (2006.01)
A61B 8/00 (2006.01)
A61B 6/00 (2006.01)
G06Q 50/22 (2018.01)
G06F 40/295 (2020.01)
G06F 18/24 (2023.01)
G06F 18/2111 (2023.01)
G06V 30/194 (2022.01)

(52) **U.S. Cl.**

CPC *G06F 3/0484* (2013.01); *G06F 9/542* (2013.01); *G06F 16/245* (2019.01); *G06F 18/214* (2023.01); *G06F 18/217* (2023.01); *G06F 18/2115* (2023.01); *G06F 18/2415* (2023.01); *G06F 18/41* (2023.01); *G06F 21/6254* (2013.01); *G06N 5/04* (2013.01); *G06N 5/045* (2013.01); *G06N 20/00*

(2019.01); *G06N 20/20* (2019.01); *G06Q 20/14* (2013.01); *G06T 3/40* (2013.01); *G06T 5/002* (2013.01); *G06T 5/008* (2013.01); *G06T 5/50* (2013.01); *G06T 7/0012* (2013.01); *G06T 7/0014* (2013.01); *G06T 7/10* (2017.01); *G06T 7/11* (2017.01); *G06T 7/187* (2017.01); *G06T 7/44* (2017.01); *G06T 7/97* (2017.01); *G06T 11/001* (2013.01); *G06T 11/006* (2013.01); *G06T 11/206* (2013.01); *G06V 10/225* (2022.01); *G06V 10/25* (2022.01); *G06V 10/764* (2022.01); *G06V 10/82* (2022.01); *G06V 30/19173* (2022.01); *G06V 40/171* (2022.01); *G16H 10/20* (2018.01); *G16H 10/60* (2018.01); *G16H 15/00* (2018.01); *G16H 30/20* (2018.01); *G16H 30/40* (2018.01); *G16H 40/20* (2018.01); *G16H 50/20* (2018.01); *H04L 67/01* (2022.05); *H04L 67/12* (2013.01); *A61B 5/055* (2013.01); *A61B 6/032* (2013.01); *A61B 6/5217* (2013.01); *A61B 8/4416* (2013.01); *G06F 18/2111* (2023.01); *G06F 18/24* (2023.01); *G06F 40/295* (2020.01); *G06Q 50/22* (2013.01); *G06T 7/70* (2017.01); *G06T 2200/24* (2013.01); *G06T 2207/10048* (2013.01); *G06T 2207/10081* (2013.01); *G06T 2207/10088* (2013.01); *G06T 2207/10116* (2013.01); *G06T 2207/10132* (2013.01); *G06T 2207/20076* (2013.01); *G06T 2207/20081* (2013.01); *G06T 2207/20084* (2013.01); *G06T 2207/30004* (2013.01); *G06T 2207/30008* (2013.01); *G06T 2207/30016* (2013.01); *G06T 2207/30061* (2013.01); *G06V 30/194* (2022.01); *G06V 2201/03* (2022.01); *G16H 50/30* (2018.01); *G16H 50/70* (2018.01)

(58) **Field of Classification Search**

CPC . *G06T 2207/10132*; *G06T 2207/30004*; *G16H 50/70*; *G16H 50/30*; *G06K 9/6229*; *G06K 9/6267*; *G06K 9/66*; *G06K 2209/05*

See application file for complete search history.

(56)

References Cited

U.S. PATENT DOCUMENTS

6,819,790 B2	11/2004	Suzuki
6,937,776 B2	8/2005	Li
7,123,762 B2	10/2006	Giger
7,418,123 B2	8/2008	Giger
7,813,822 B1	10/2010	Hoffberg
8,121,362 B2	2/2012	Zhan
8,303,505 B2	11/2012	Webler
8,331,637 B2	12/2012	Bar-Aviv
8,379,950 B2	2/2013	Ye
8,600,133 B2	12/2013	Buelow
8,885,898 B2	11/2014	Liu
9,165,360 B1	10/2015	Bates
9,349,178 B1	5/2016	Itu
9,501,863 B1	11/2016	Fram
9,569,736 B1	2/2017	Ghesu
9,579,518 B2	2/2017	Gertner
9,760,978 B1	9/2017	Lu
10,140,421 B1	11/2018	Bernard
10,304,198 B2	5/2019	Yan
10,340,041 B2	7/2019	Chan
2002/0186818 A1	12/2002	Arnaud
2003/0016850 A1	1/2003	Kaufman
2004/0064029 A1	4/2004	Summers
2004/0147840 A1	7/2004	Duggirala
2004/0181431 A1	9/2004	Kuth

(56)

References Cited

OTHER PUBLICATIONS

U.S. PATENT DOCUMENTS

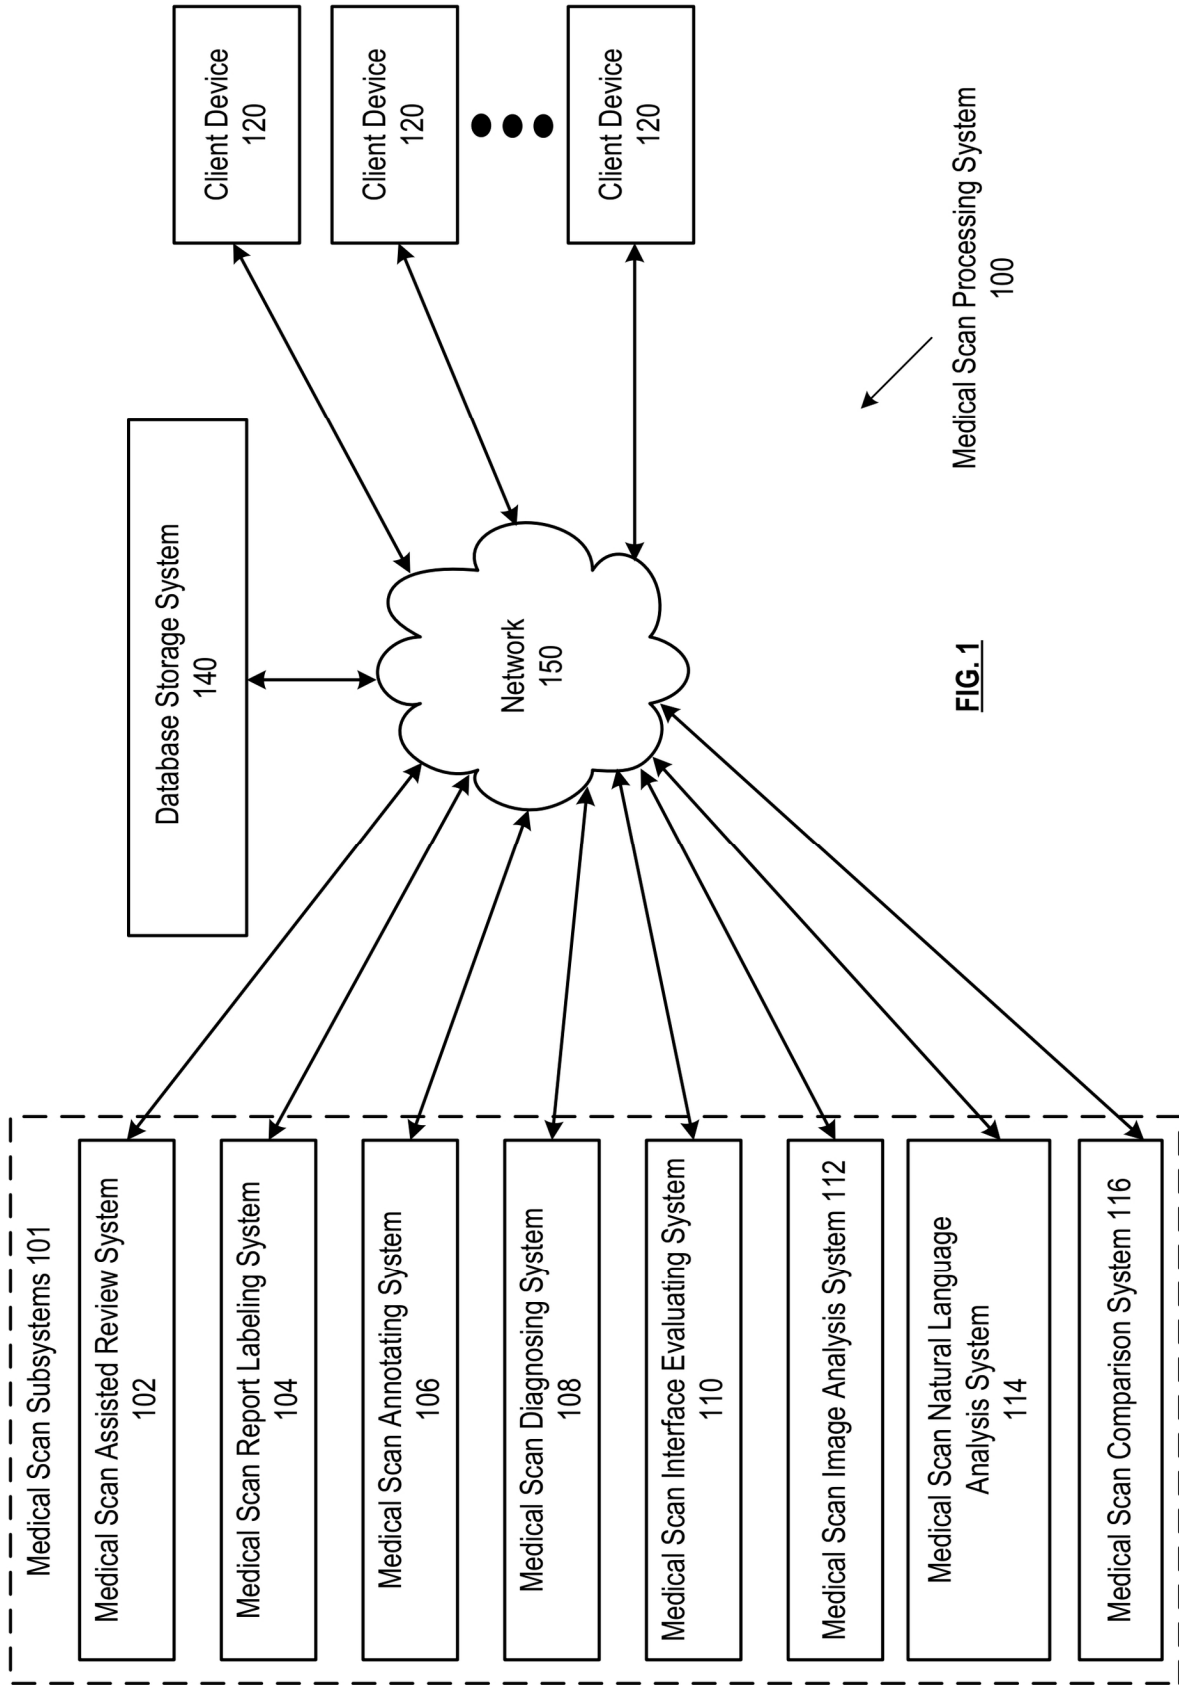
2004/0252870 A1 12/2004 Reeves
 2005/0027566 A1 2/2005 Haskell
 2005/0207630 A1 9/2005 Chan
 2005/0283450 A1 12/2005 Matsugu
 2007/0004980 A1 1/2007 Warner
 2008/0004505 A1 1/2008 Kapit
 2008/0015418 A1 1/2008 Jarrell
 2008/0021834 A1 1/2008 Holla
 2008/0205717 A1 8/2008 Reeves
 2008/0243539 A1 10/2008 Barish et al.
 2008/0267483 A1 10/2008 Zhan
 2009/0177495 A1 7/2009 Abousy
 2009/0222388 A1 9/2009 Hua
 2009/0228299 A1 9/2009 Kangarloo
 2010/0088117 A1 4/2010 Belden
 2010/0278405 A1 11/2010 Kakadiaris
 2012/0130734 A1 5/2012 White
 2013/0018674 A1* 1/2013 Bedi G06Q 10/00
 705/3
 2013/0124527 A1 5/2013 Lee
 2014/0219500 A1 8/2014 Moehrl
 2014/0328517 A1 11/2014 Gluncic
 2014/0341471 A1 11/2014 Ono et al.
 2015/0031979 A1 1/2015 Rappaport et al.
 2015/0063667 A1 3/2015 Sprencz
 2015/0230773 A1 8/2015 Cho
 2015/0254843 A1 9/2015 Brown
 2015/0305706 A1 10/2015 Kanik
 2016/0019695 A1 1/2016 Chukka
 2016/0027175 A1 1/2016 Kim et al.
 2016/0048972 A1 2/2016 Kam
 2016/0104281 A1 4/2016 Grady
 2016/0174902 A1 6/2016 Georgescu
 2016/0203281 A1 7/2016 Zalis
 2016/0239962 A1* 8/2016 Markin G16H 40/20
 2016/0260211 A1 9/2016 Gillies
 2016/0314588 A1 10/2016 Harper
 2016/0343127 A1 11/2016 Miller
 2016/0350919 A1 12/2016 Steigauf
 2016/0364862 A1 12/2016 Reicher
 2017/0004619 A1 1/2017 Liang
 2017/0024517 A1 1/2017 Biegert
 2017/0053074 A1 2/2017 Enzmann
 2017/0068780 A1 3/2017 Dobrean
 2017/0116497 A1 4/2017 Georgescu
 2017/0330320 A1 11/2017 Lynch
 2017/0337343 A1 11/2017 Kakadiaris
 2018/0025255 A1 1/2018 Poole
 2018/0033144 A1 2/2018 Risman
 2018/0060535 A1 3/2018 Reicher
 2018/0060691 A1 3/2018 Bernal
 2018/0114595 A1 4/2018 Stern
 2018/0137244 A1* 5/2018 Sorenson G16H 30/20
 2018/0204111 A1 7/2018 Zadeh
 2018/0260950 A1 9/2018 Samset
 2018/0268737 A1 9/2018 Garnavi
 2018/0315182 A1 11/2018 Rapaka
 2018/0342055 A1 11/2018 Lyman
 2019/0142519 A1 5/2019 Siemionow
 2020/0288972 A1* 9/2020 Park G06N 99/00

FOREIGN PATENT DOCUMENTS

CN 106551704 A 4/2017
 EP 1503359 A1 2/2005
 GB 2457022 A 8/2009
 KR 1020100014065 A 2/2010

Andersch, Michael; Inference: The Next Step in GPU-Accelerated Deep Learning; <https://devblogs.nvidia.com/parallelforall/inference-next-step-gpu-accelerated-deep-learning/>; Nov. 11, 2015; 7 pages.
 Armato et al., "Assessment of Radiologist Performance in the Detecting of Lung Nodules: Dependence on the Definition of "Truth"", Jan. 2009, Academic Radiology, vol. 16, No. 1, pp. 28-38.
 Attaluri, et al.; Efficient and Accurate Abnormality Mining from Radiology Reports with Customized False Positive Reduction: arXiv:1810.000967v1; Oct. 1, 2018.
 Choi, Y., Chiu, C. Y. I., & Sontag, D. (2016). Learning low-dimensional representations of medical concepts. AM IA Summits on Translational Science Proceedings, 2016, 41. (Year: 2016).
 International Searching Authority; International Search Report and Written Opinion; International Application No. PCT/US2018/032927; dated Sep. 14, 2018; 9 pgs.
 Jamie A. Weydert, MD, Barry R. De Young, MD and Michael B. Cohen, MD, A preliminary Diagnosis Service Provides Prospective Blinded Dual-Review of All General Surgical Pathology Cases in an Academic Practice, 2005, Am J Surg Pathol, vol. 29, pp. 801-805 (Year: 2005).
 Long, Jonathan, Evan Shelhamer, and Trevor Darrell. "Fully convolutional networks for semantic segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2015. (Year: 2015).
 Minnaar, Alex; Deep Learning Basics: Neural Networks, Backpropagation and Stochastic Gradient Descent; <http://alexminnaar.com/deep-learning-basics-neural-networks-backpropagation-and-stochastic-gradient-descent.html>; Feb. 14, 2015; 11 pages.
 Olah, Christopher; Calculus on Computational Graphs: Backpropagation; <http://colah.github.io/posts/2015-08-Backprop/>; Aug. 31, 2015; 7 pages.
 Pham, A.-D., Neveol, A., Lavergne, T., Yasunaga, D., Clement, O., Meyer, G., . . . Burgun, A. (2014). Natural language processing of radiology reports for the detection of thromboembolic diseases and clinically relevant incidental findings. BMC Bioinformatics, 15(1). doi: 10.1186/1471-2105-15-266 (Year: 2014).
 Pons, E., Braun, L. M. M., Hunink, M. G. M., & Kors, J. A. (2016). Natural Language Processing in Radiology: A Systematic Review. Radiology, 279(2), 329-343. doi: 10.1148/radiol.16142770.
 Pre Conference Proceedings of the 7th MICCAI BraTS Challenge (2018); BraTS Multimodal Brain Tumor Segmentation Challenge; Granada, Spain; Sep. 16, 2018; 578 pages.
 Reid, Stuart; 10 misconceptions about Neural Networks; <http://www.turingfinance.com/misconceptions-about-neural-networks/>; May 8, 2014; 24 pages.
 S. Ted Sandler, Regularized Learning with Feature Networks (Year: 2010).
 What Not to Include in a Medical Record. (Jun. 6, 2012). Retrieved from <https://www.texmed.org/Template.aspx?id=1741>. (Year: 2012).
 Wikipedia: Backpropagation; https://en.wikipedia.org/wiki/Backpropagation#Assumptions_about_the_loss_function; downloaded from the internet on 18/15/18; 12 pages.
 Wikipedia: Convolutional neural network; https://en.wikipedia.org/wiki/Convolutional_neural_network#Pooling_layer; downloaded from the internet on Jan. 15, 2018; 21 pages.
 Yao, Li, et al. "Weakly supervised medical diagnosis and localization from multiple resolutions." arXiv preprint arXiv: 1803.07703 (2018). (Year: 2018).
 Yoo Jae Chern, Remote medical-diagnosis system and method. 2010. Translated by IP.com (Year 2010).

* cited by examiner



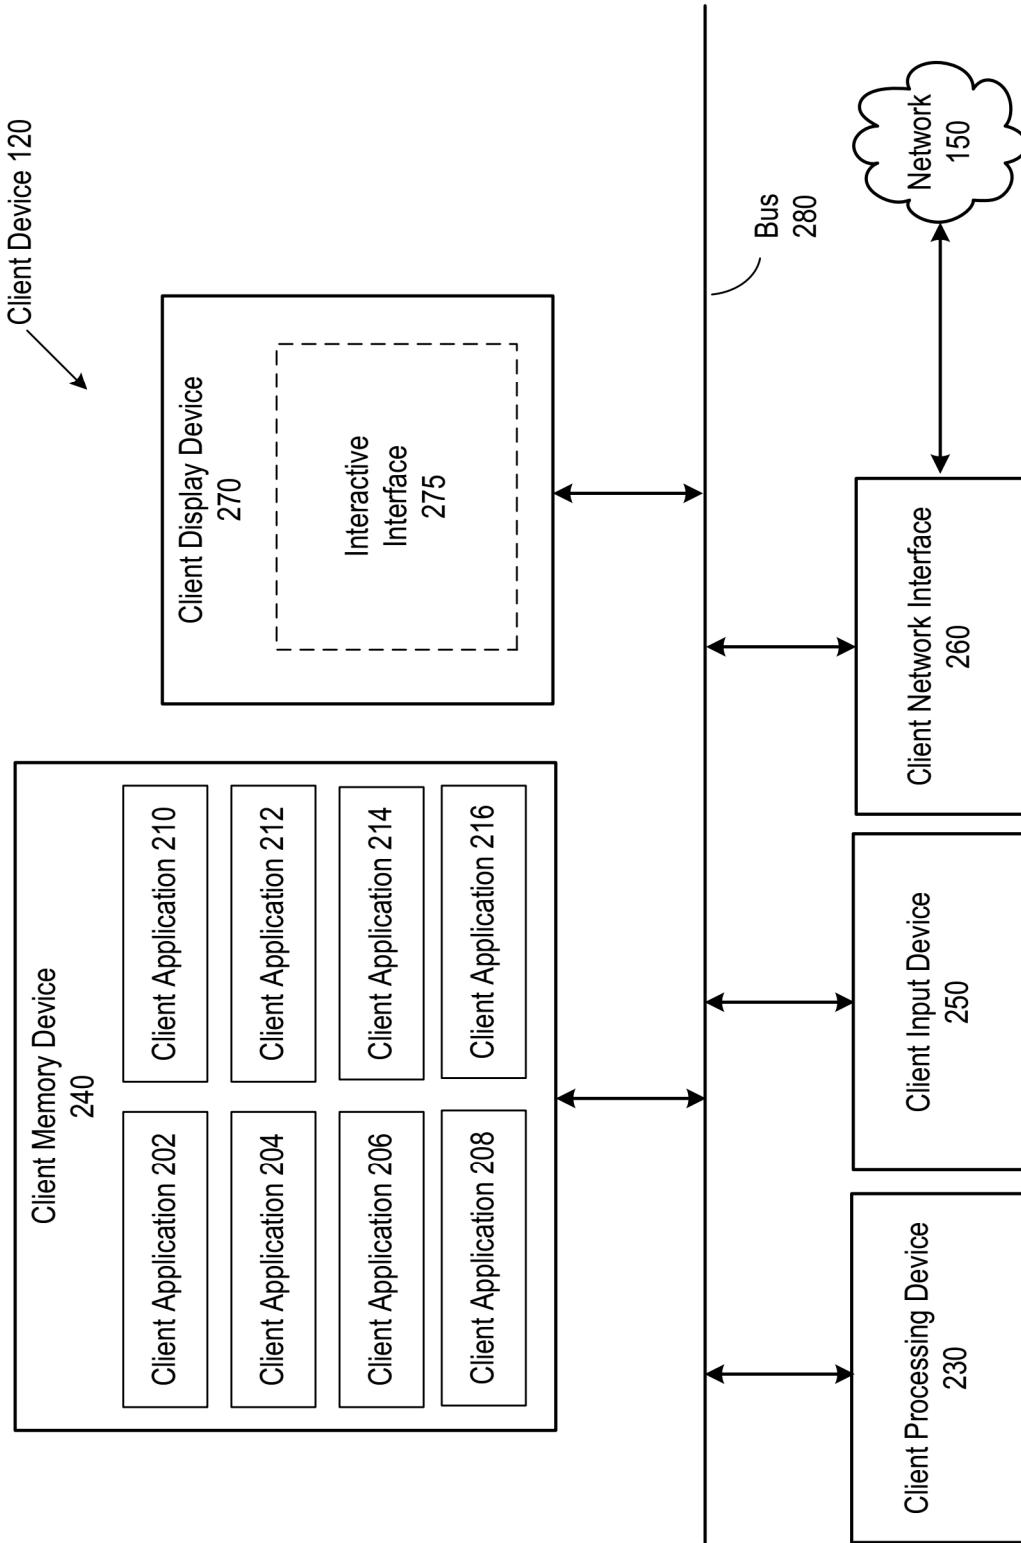


FIG. 2A




FIG. 2B




FIG. 3




FIG. 4A

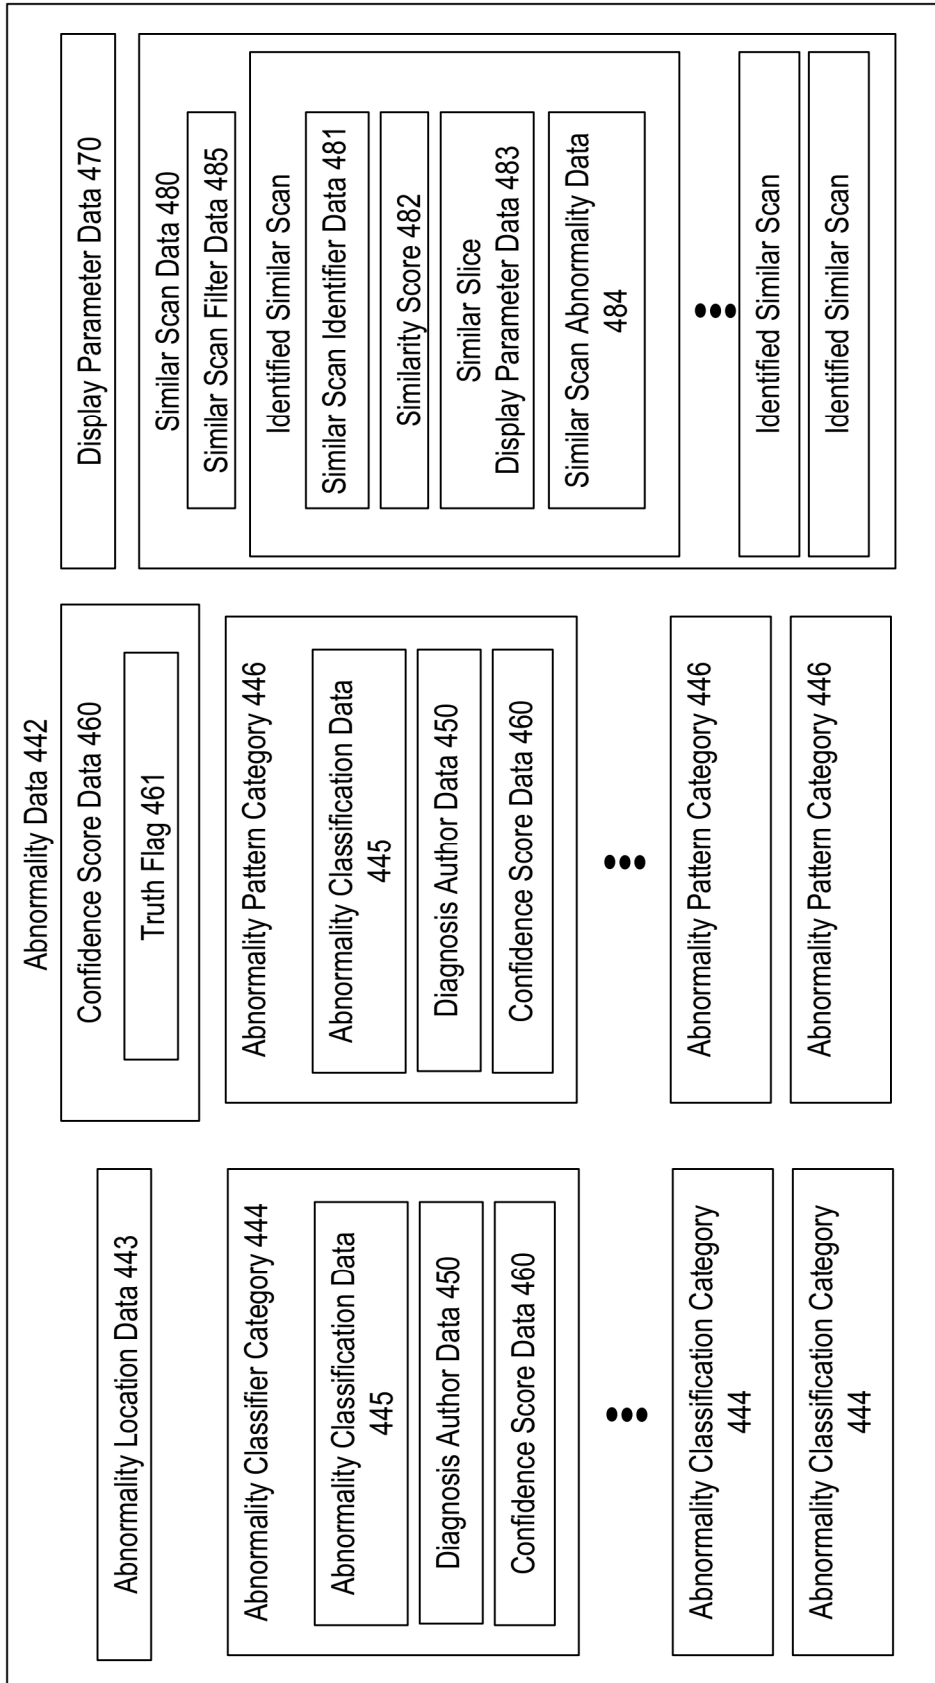


FIG. 4B

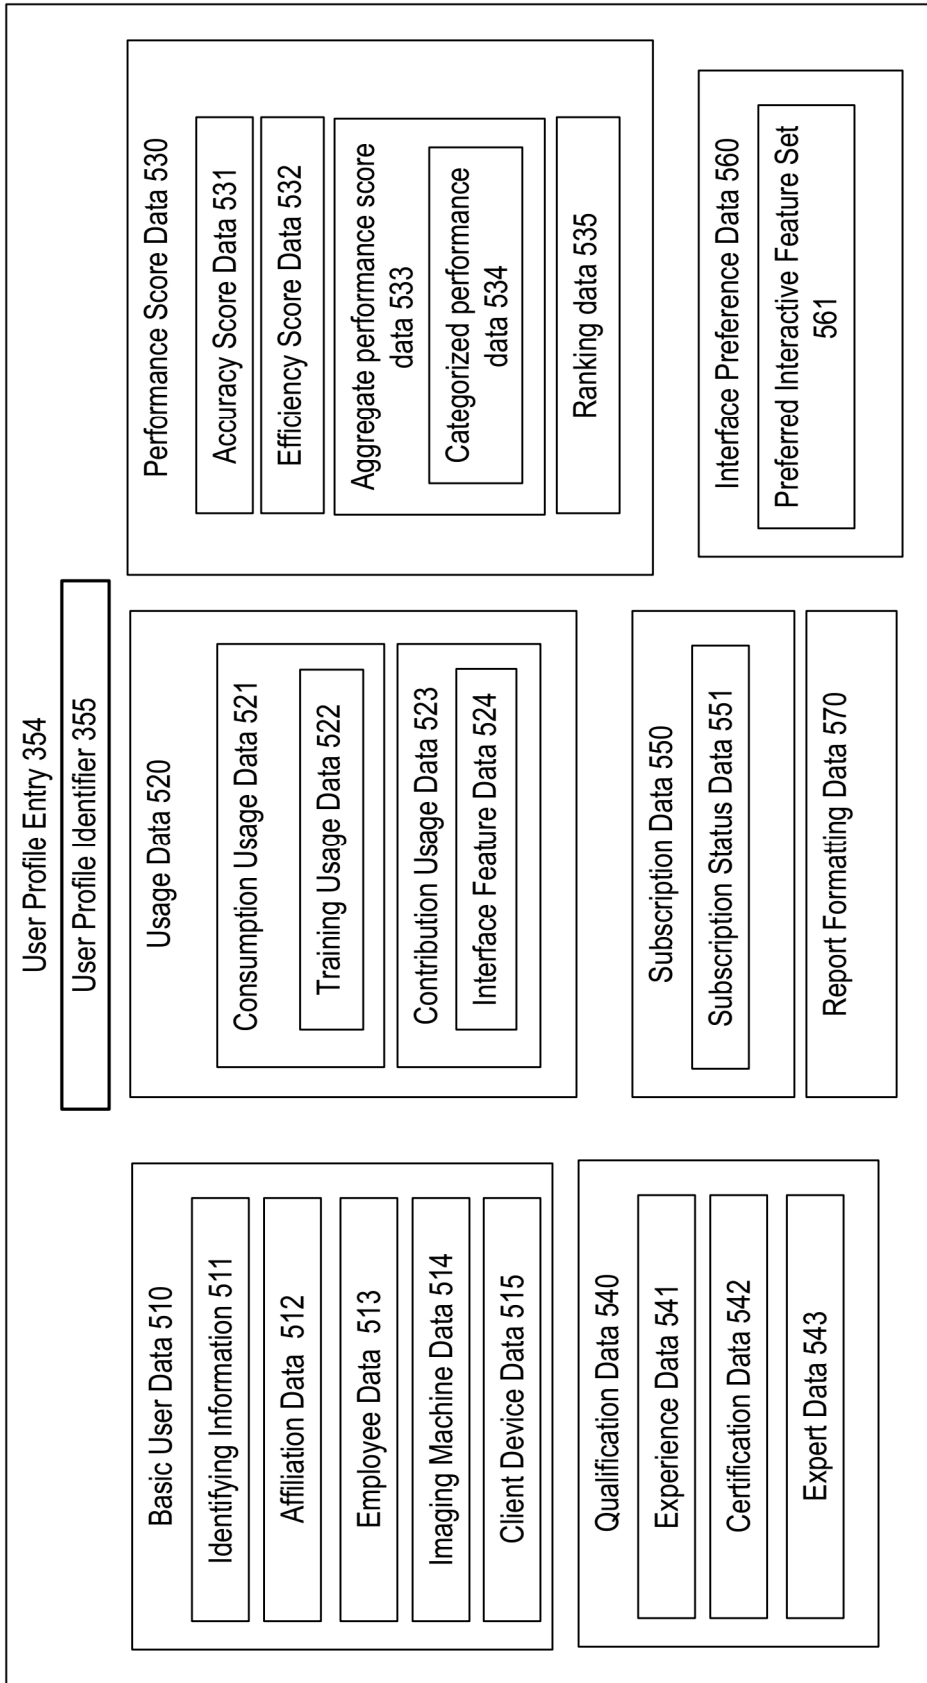


FIG. 5A

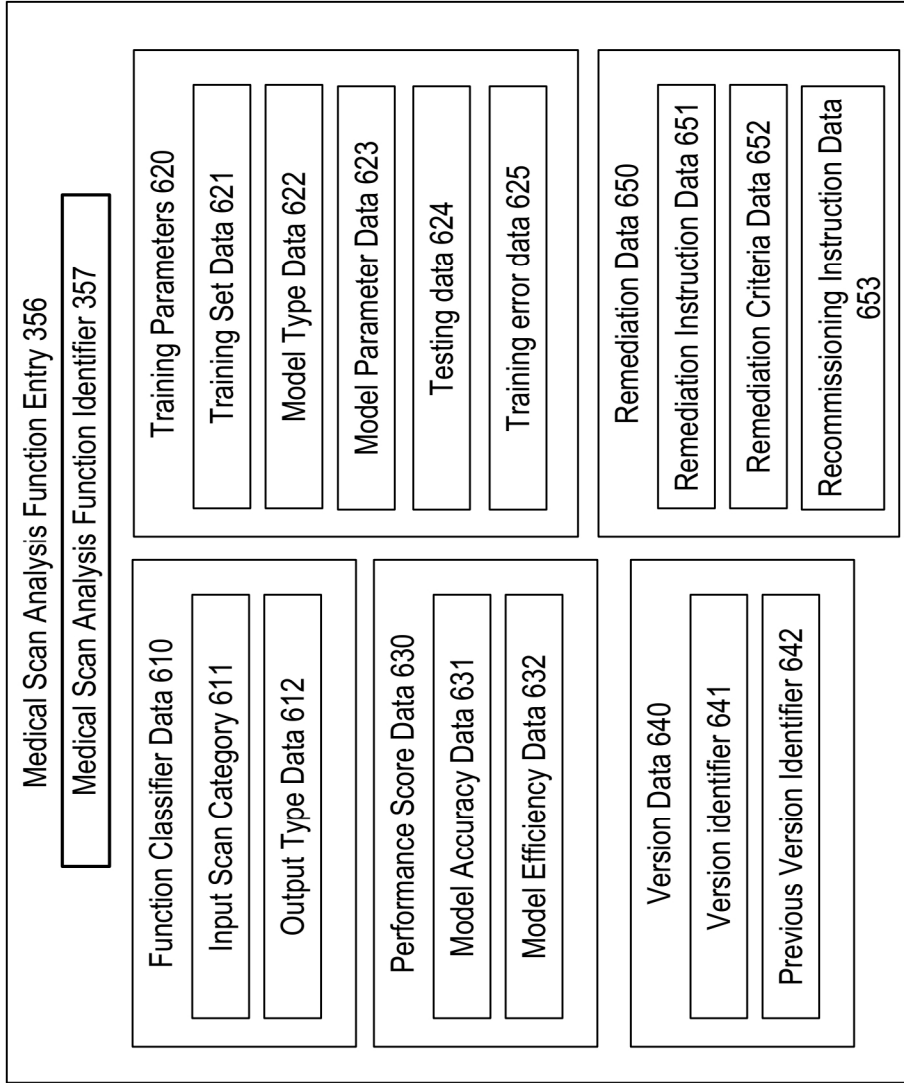


FIG. 5B

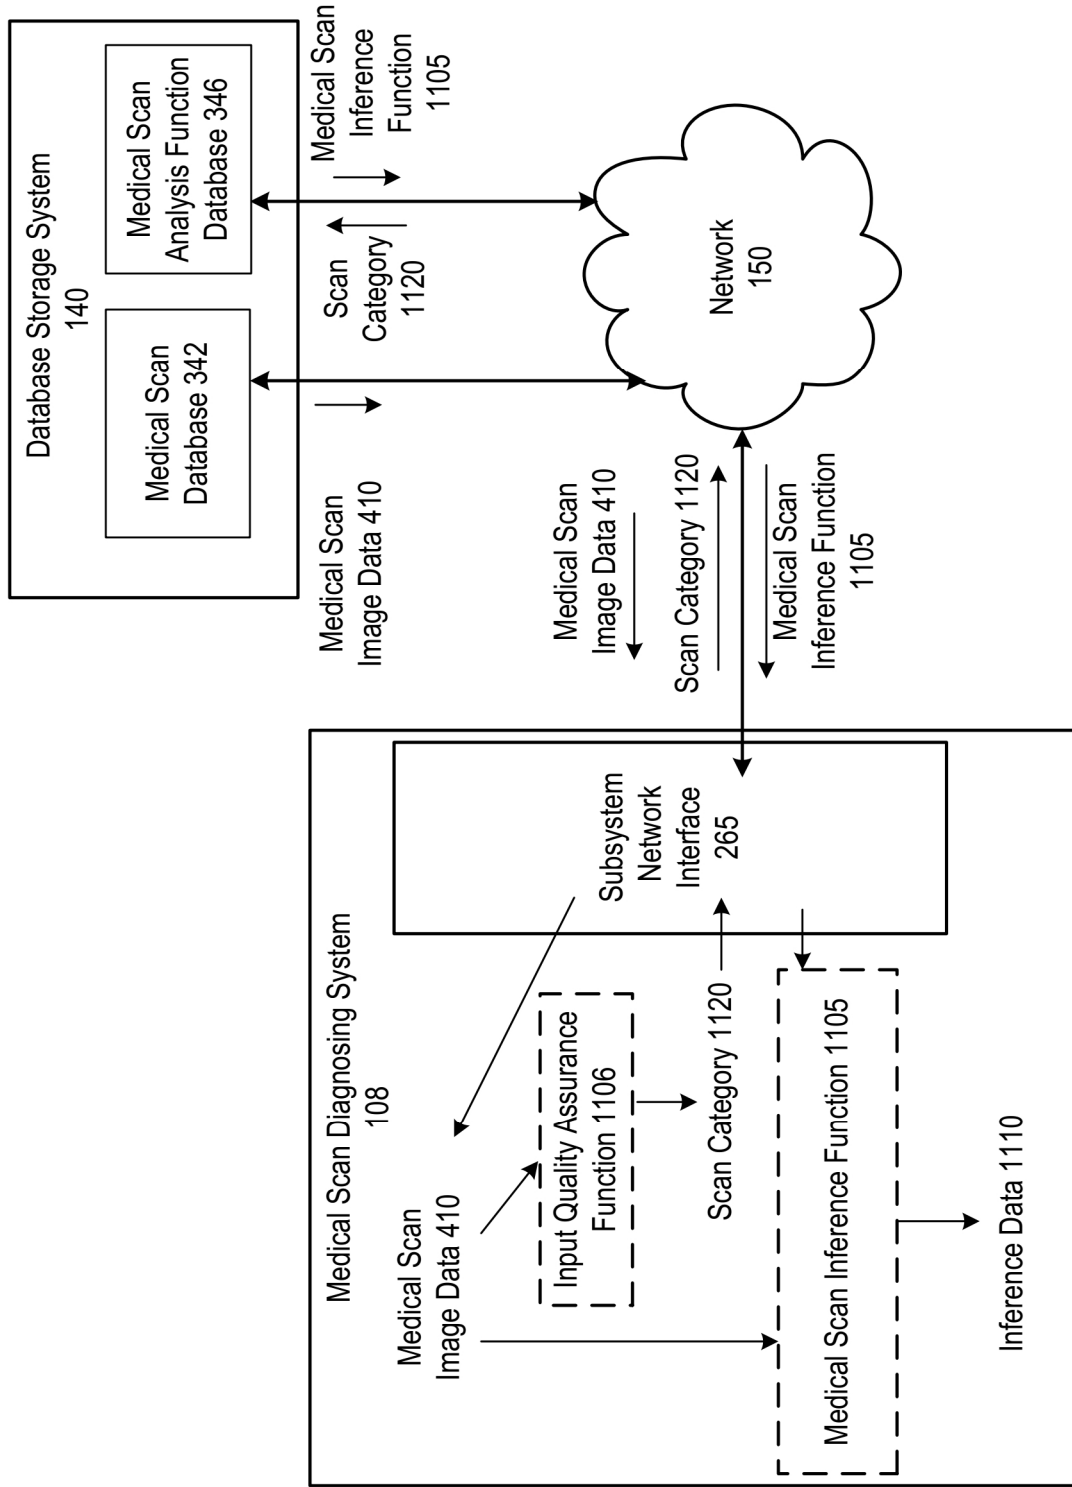


FIG. 6A

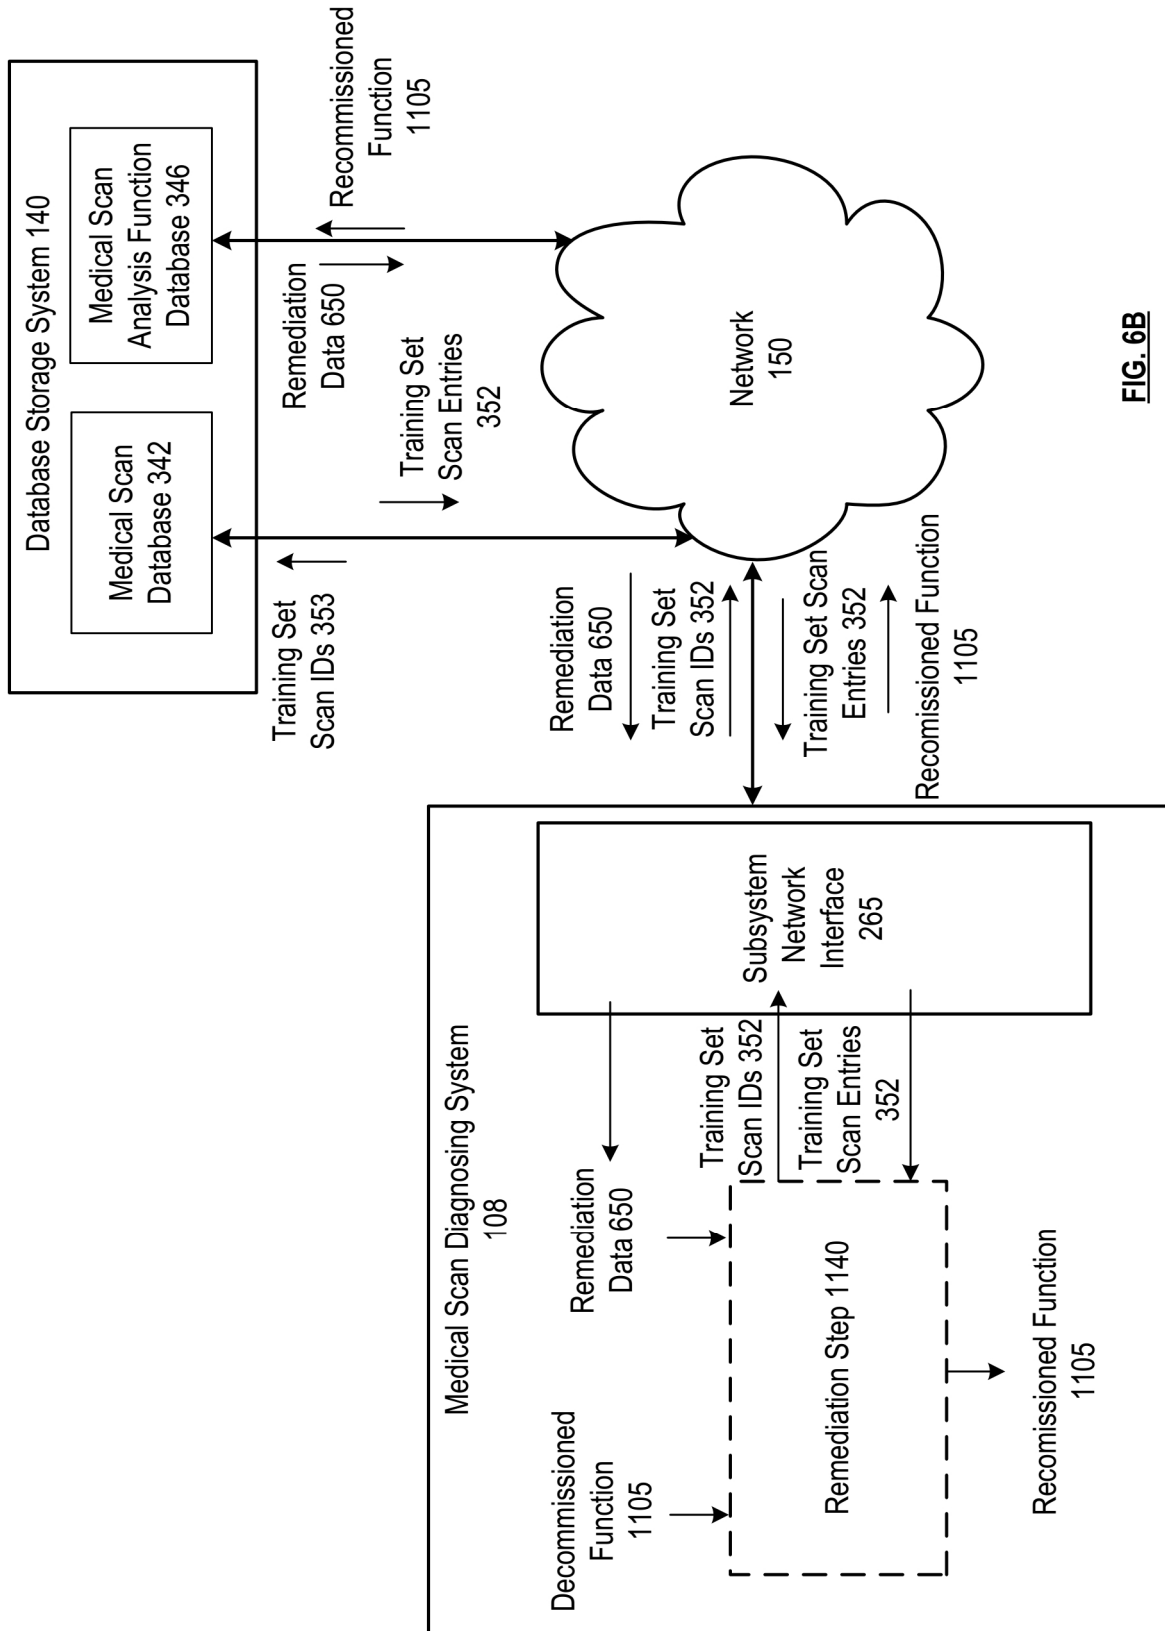


FIG. 6B




FIG. 7A




FIG.7B

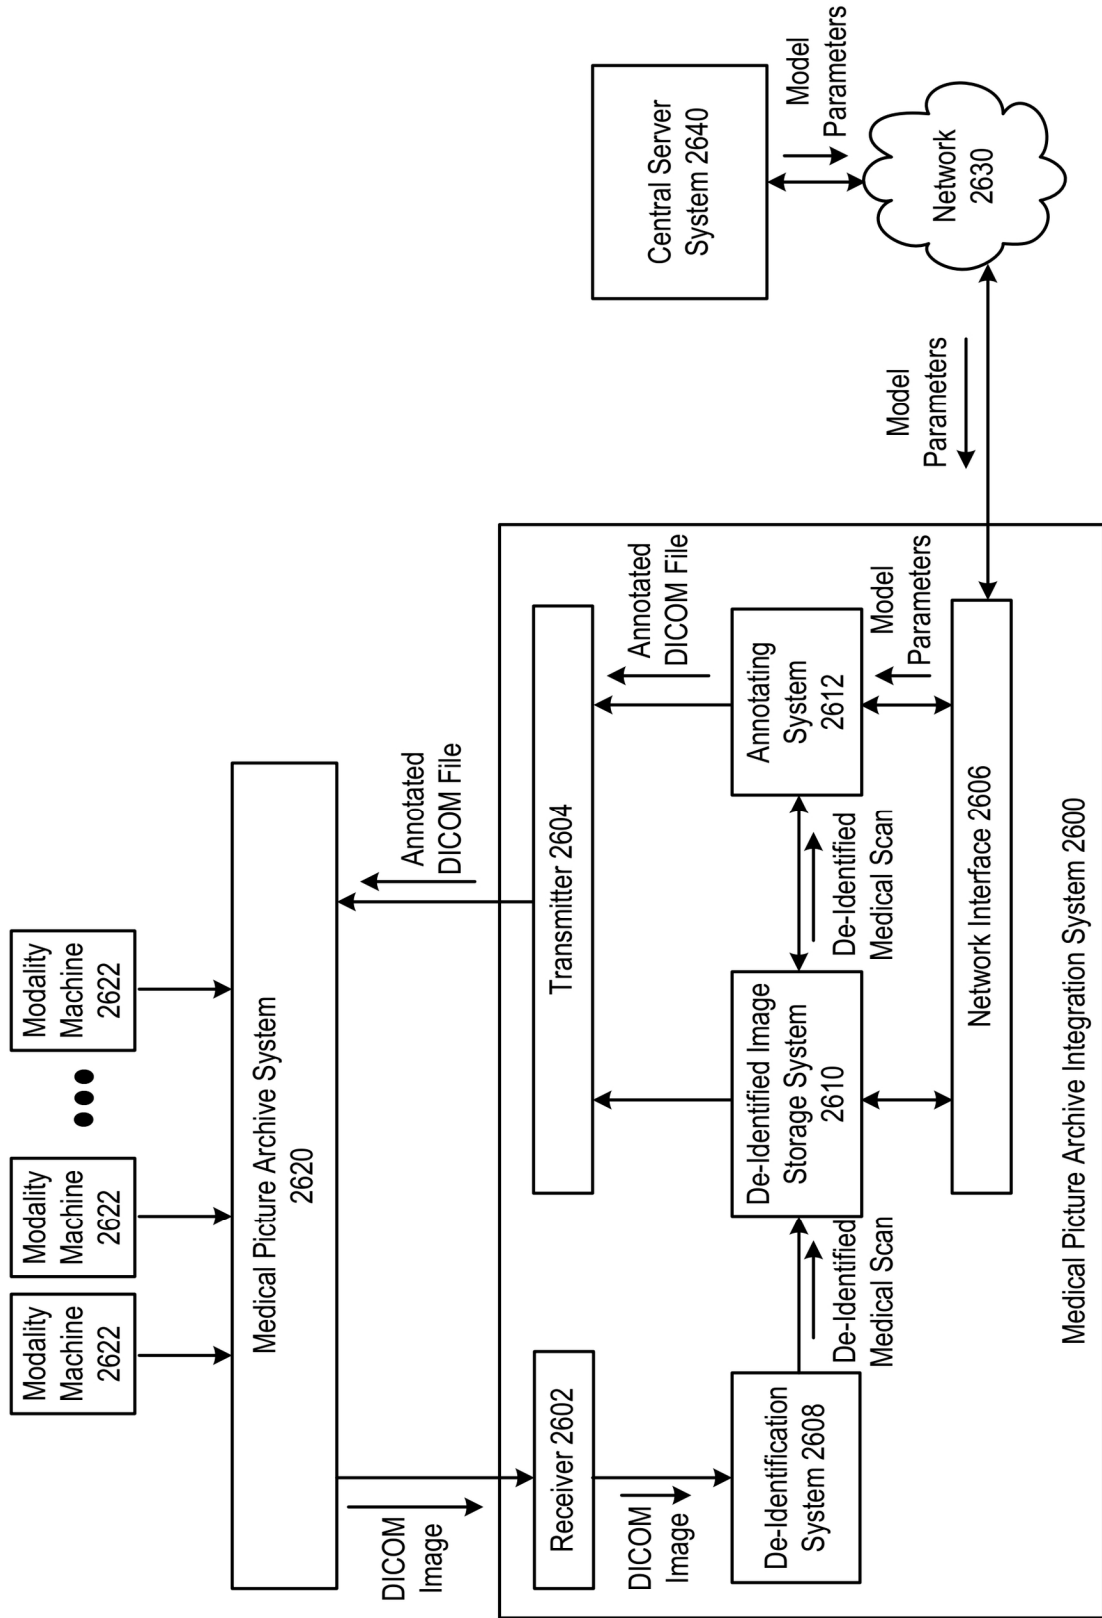


FIG. 8A




FIG. 8B

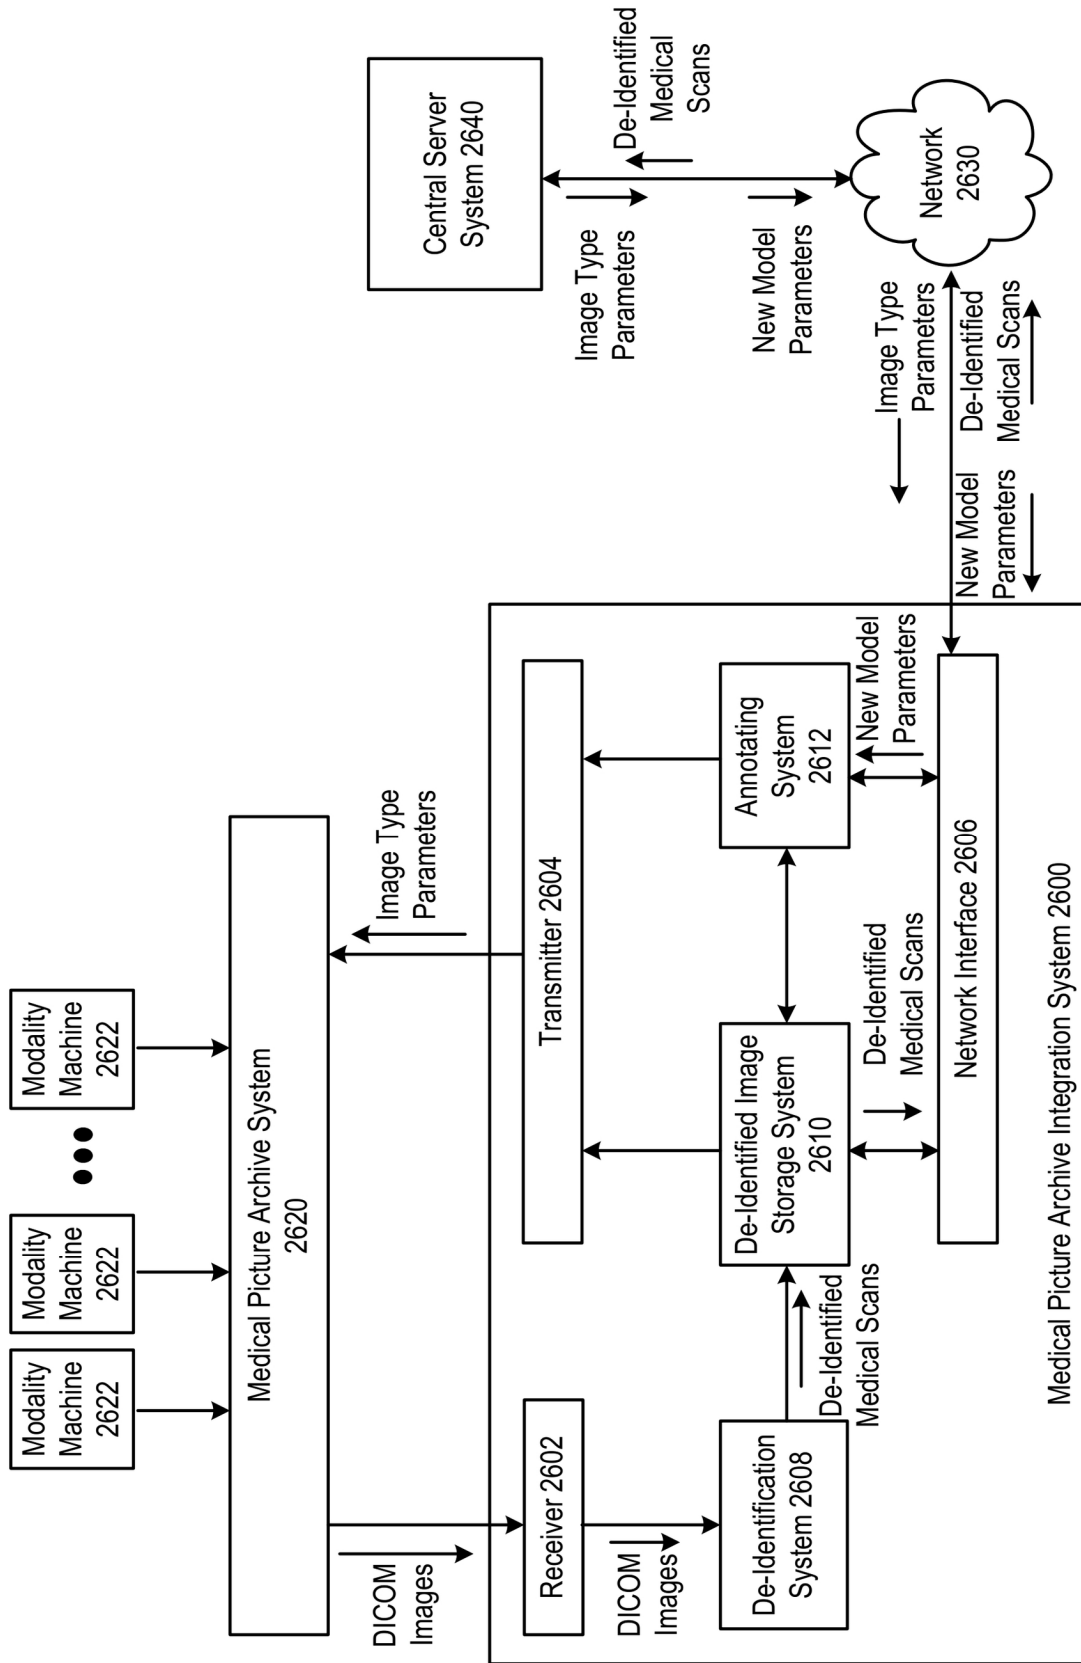


FIG. 8C

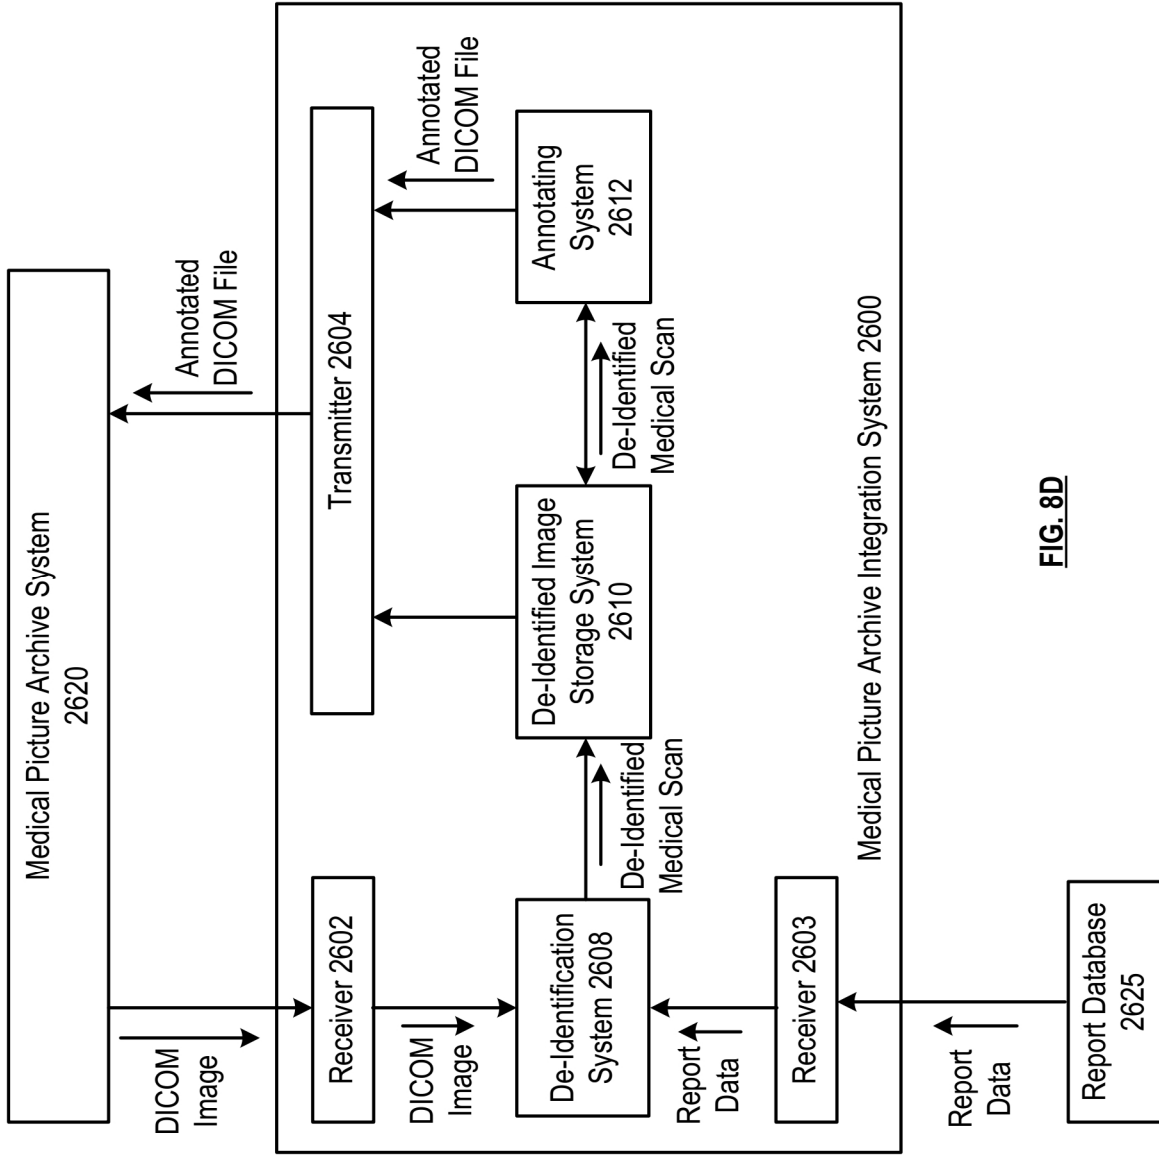


FIG. 8D

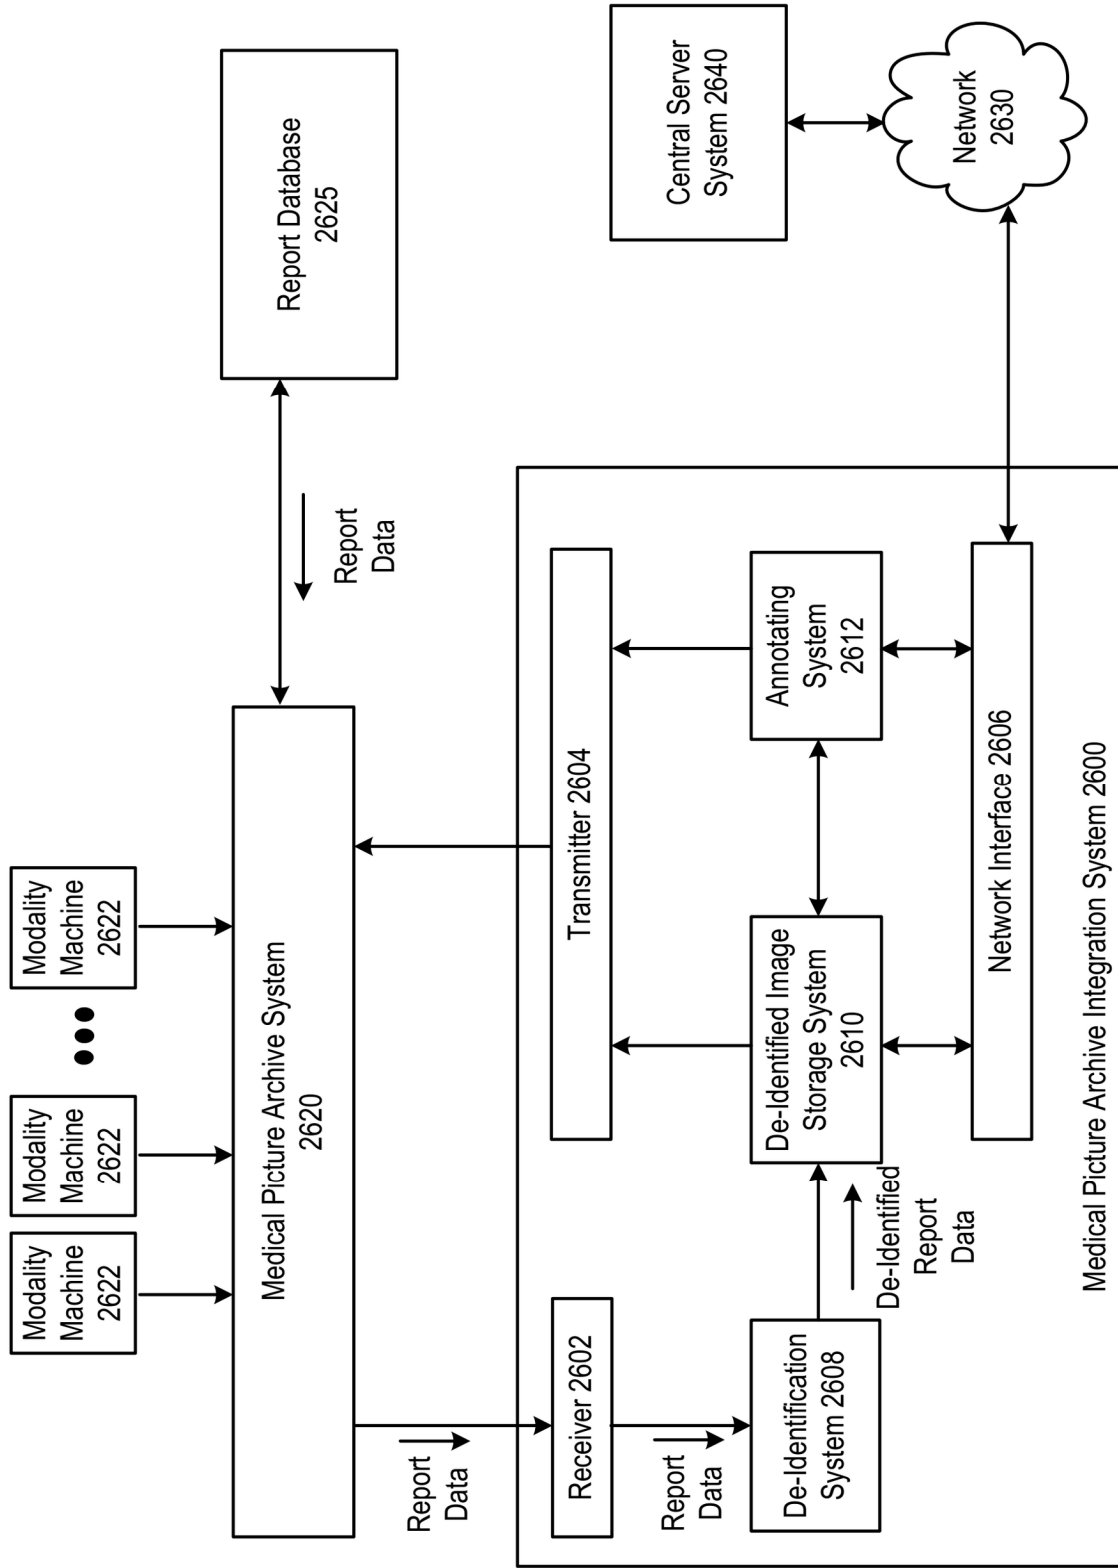


FIG. 8E

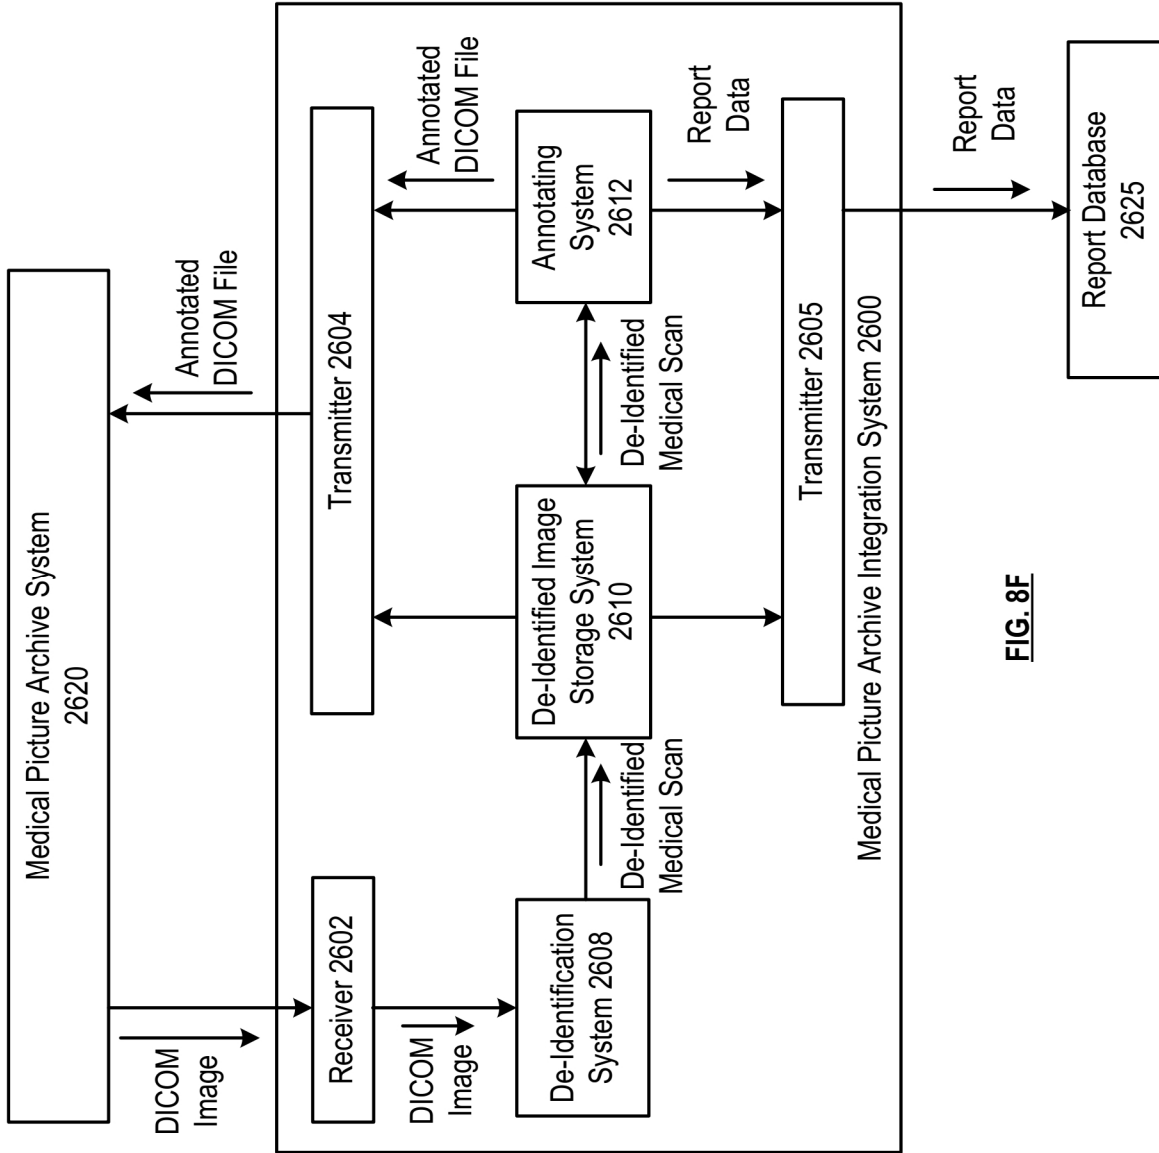


FIG. 8F

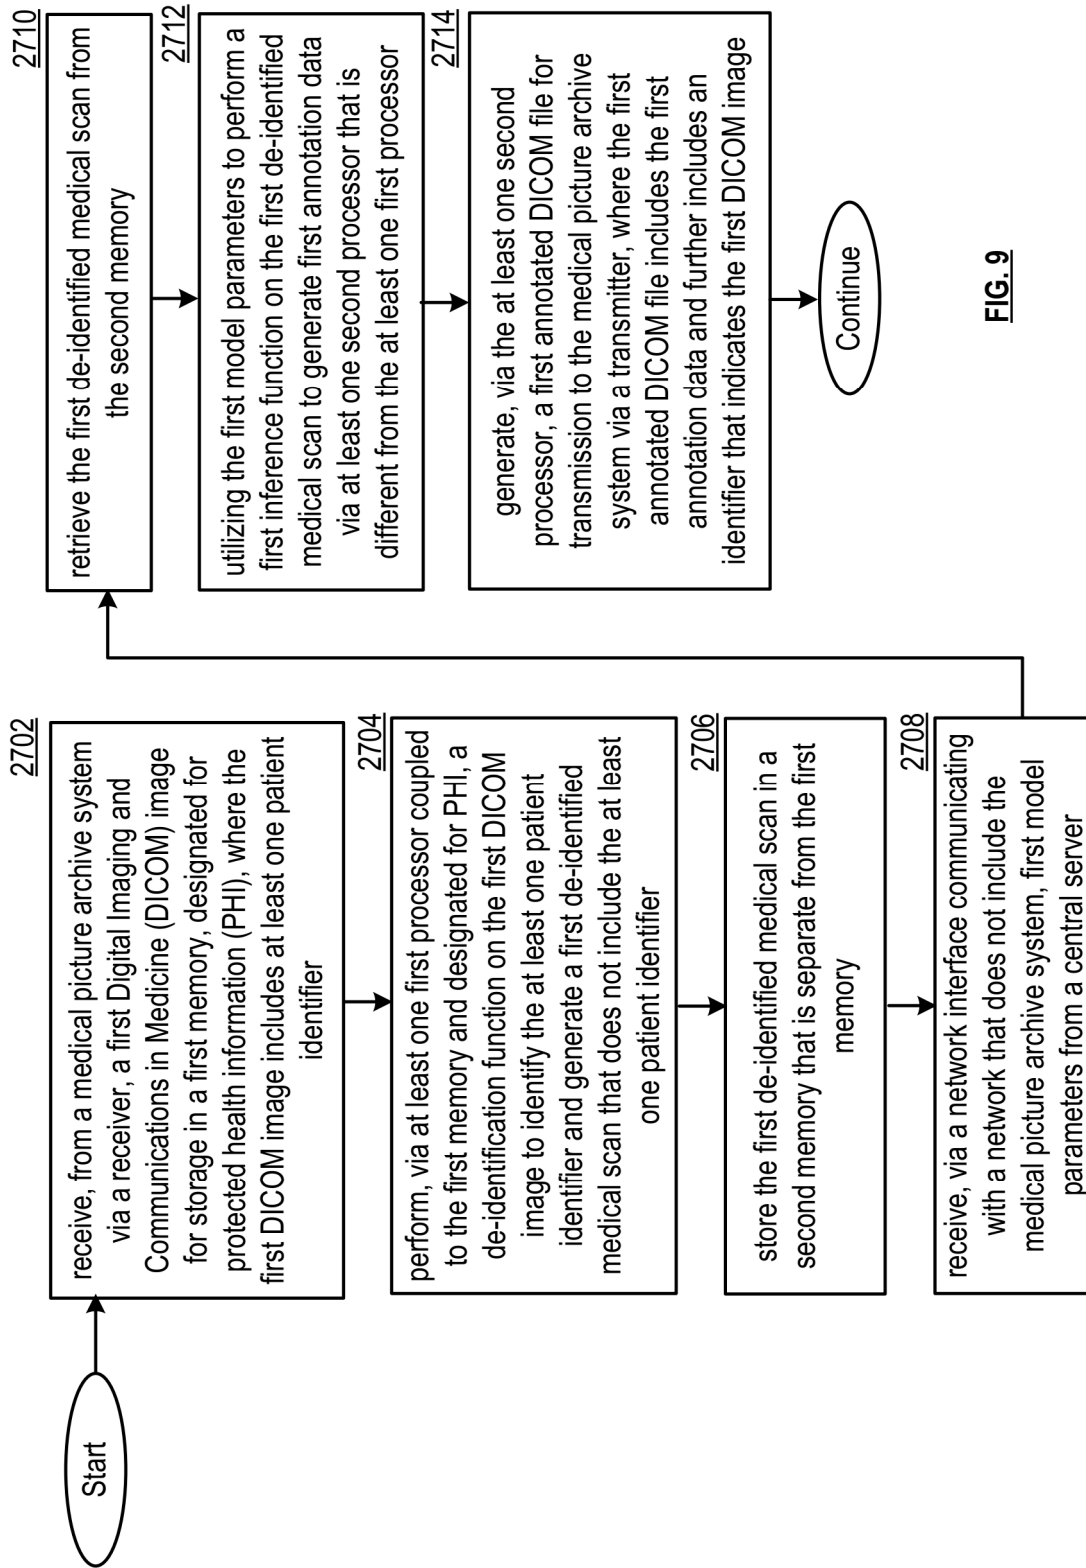


FIG. 9

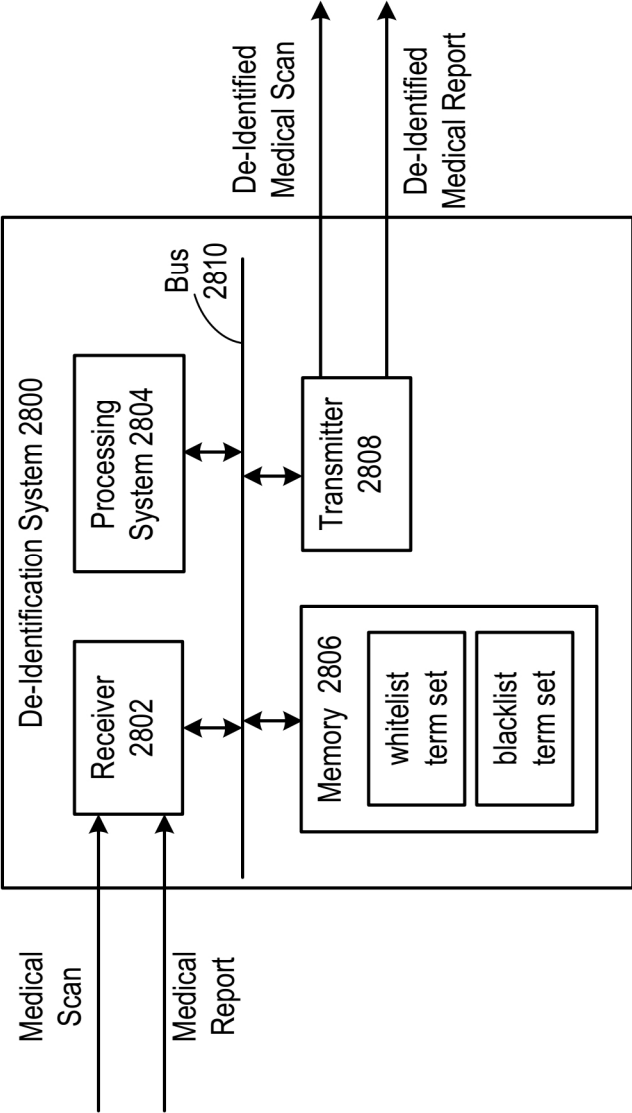


FIG. 10A




FIG. 10B

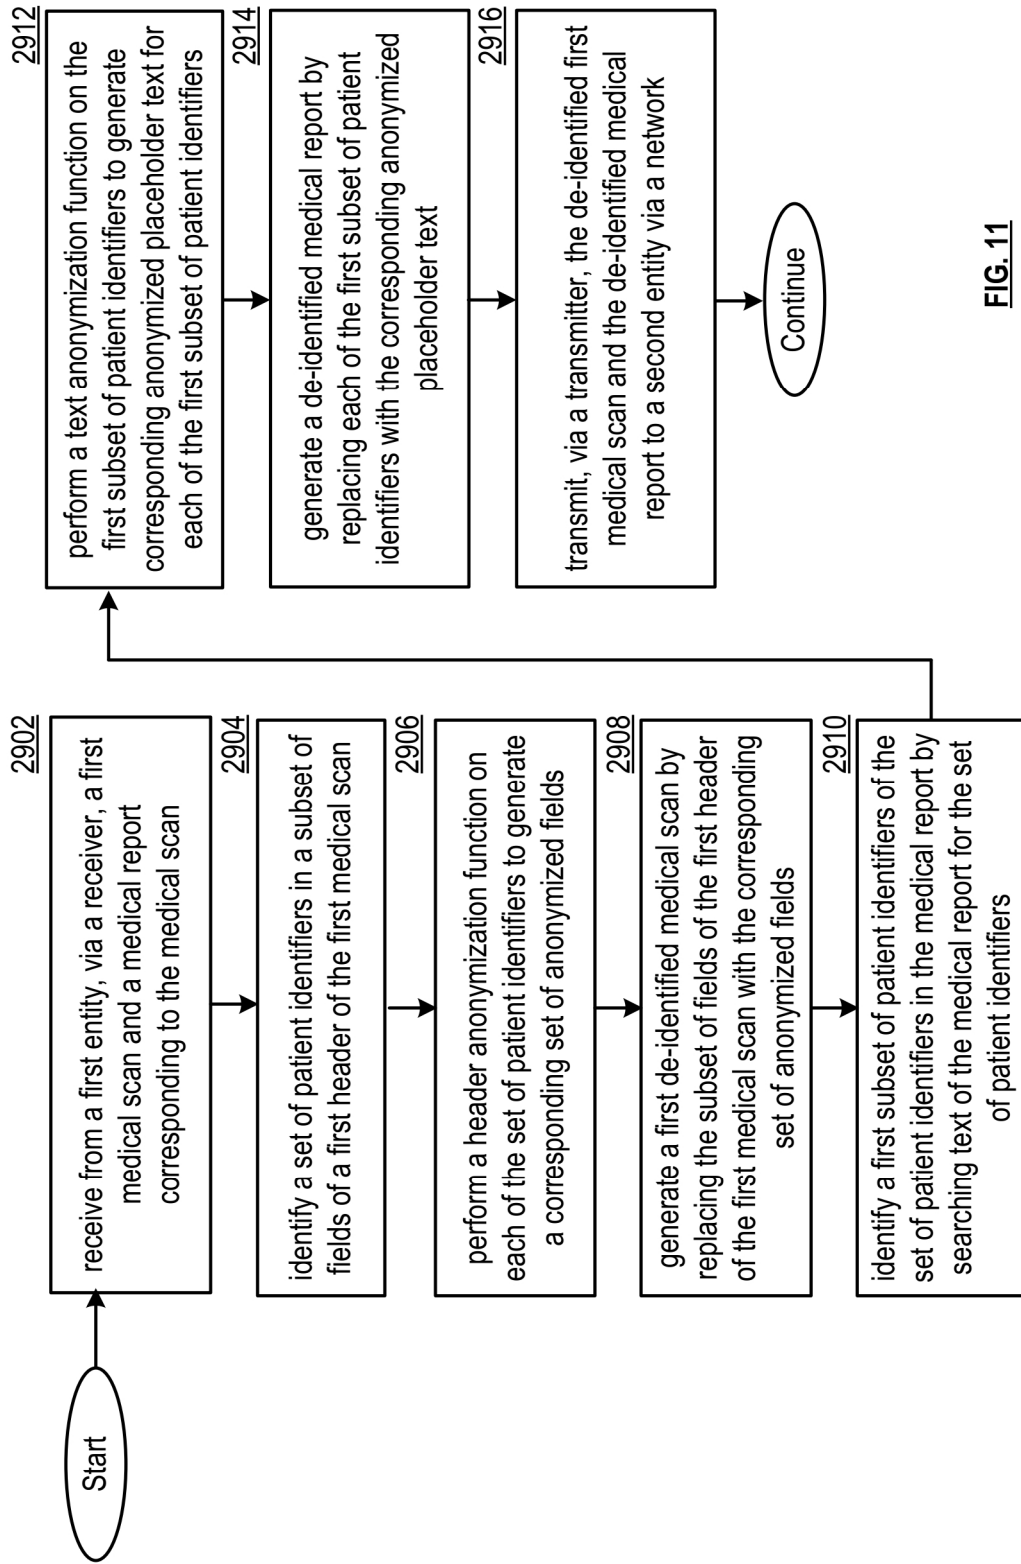


FIG. 11




FIG. 12A




FIG. 12B

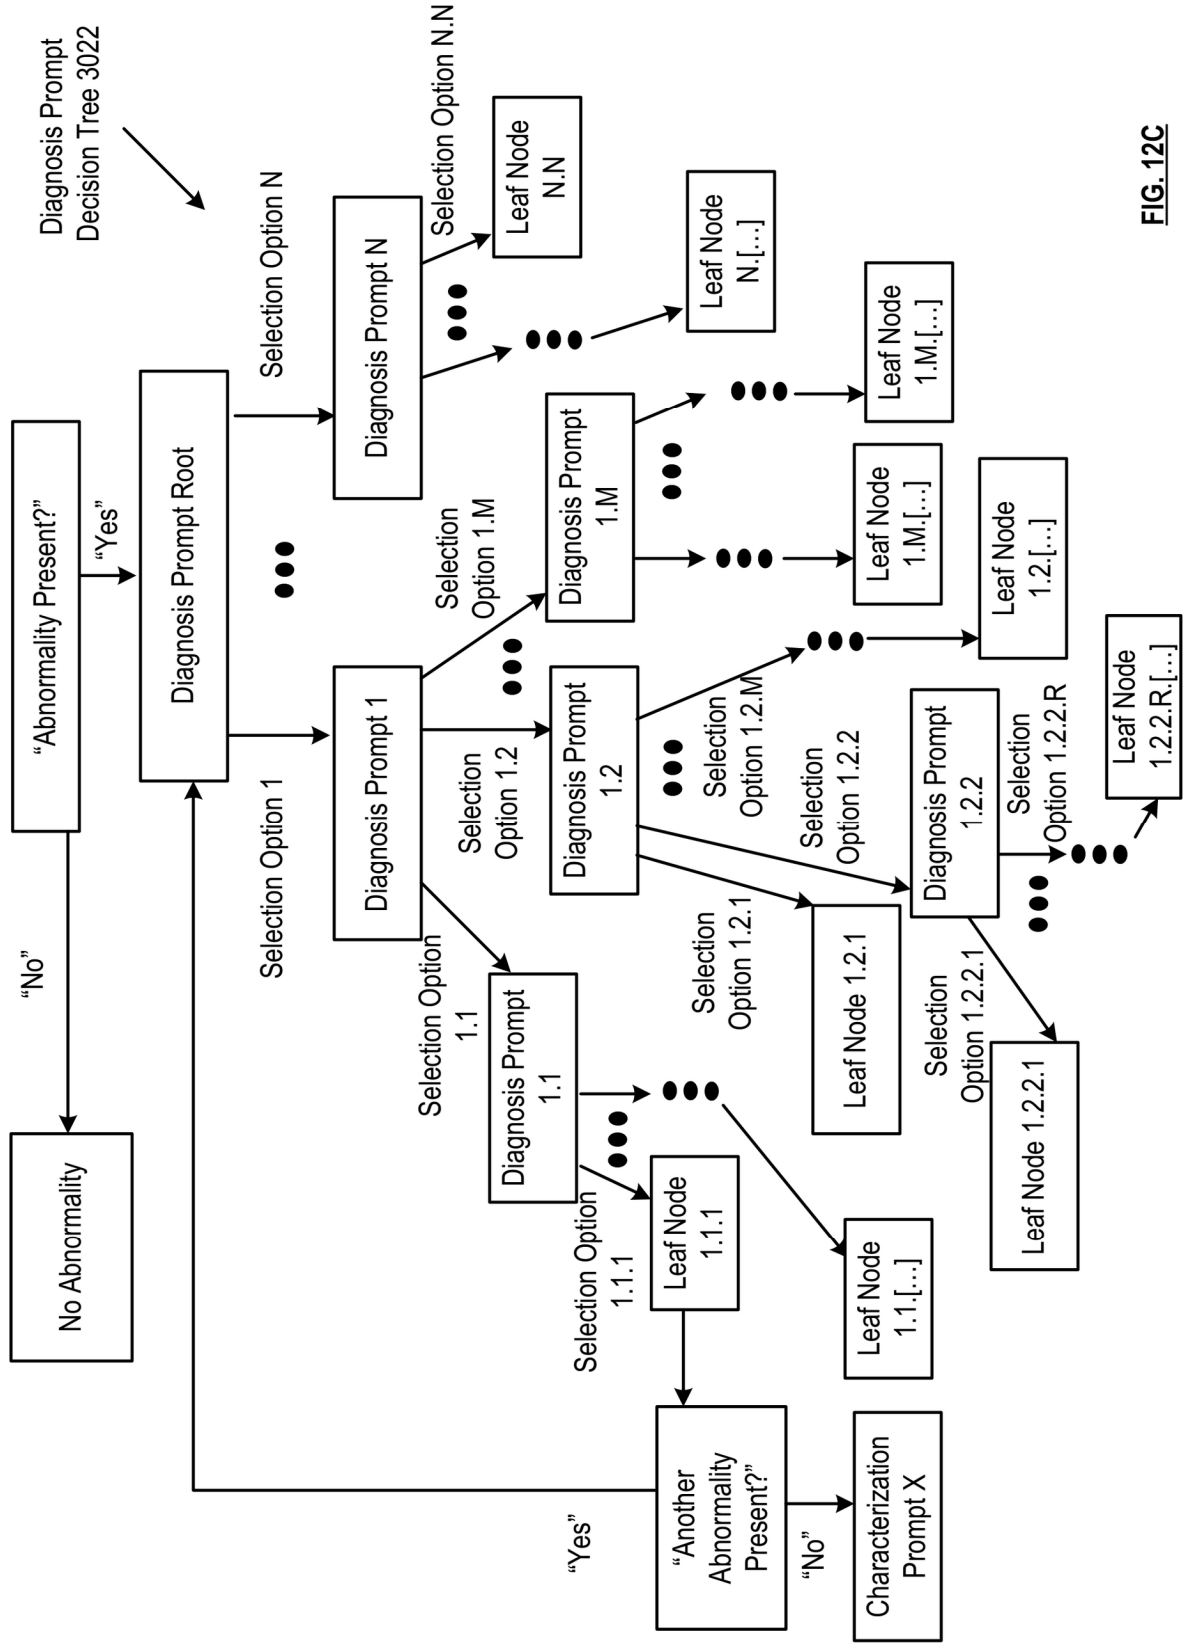


FIG. 12C




FIG. 12D




FIG. 12F

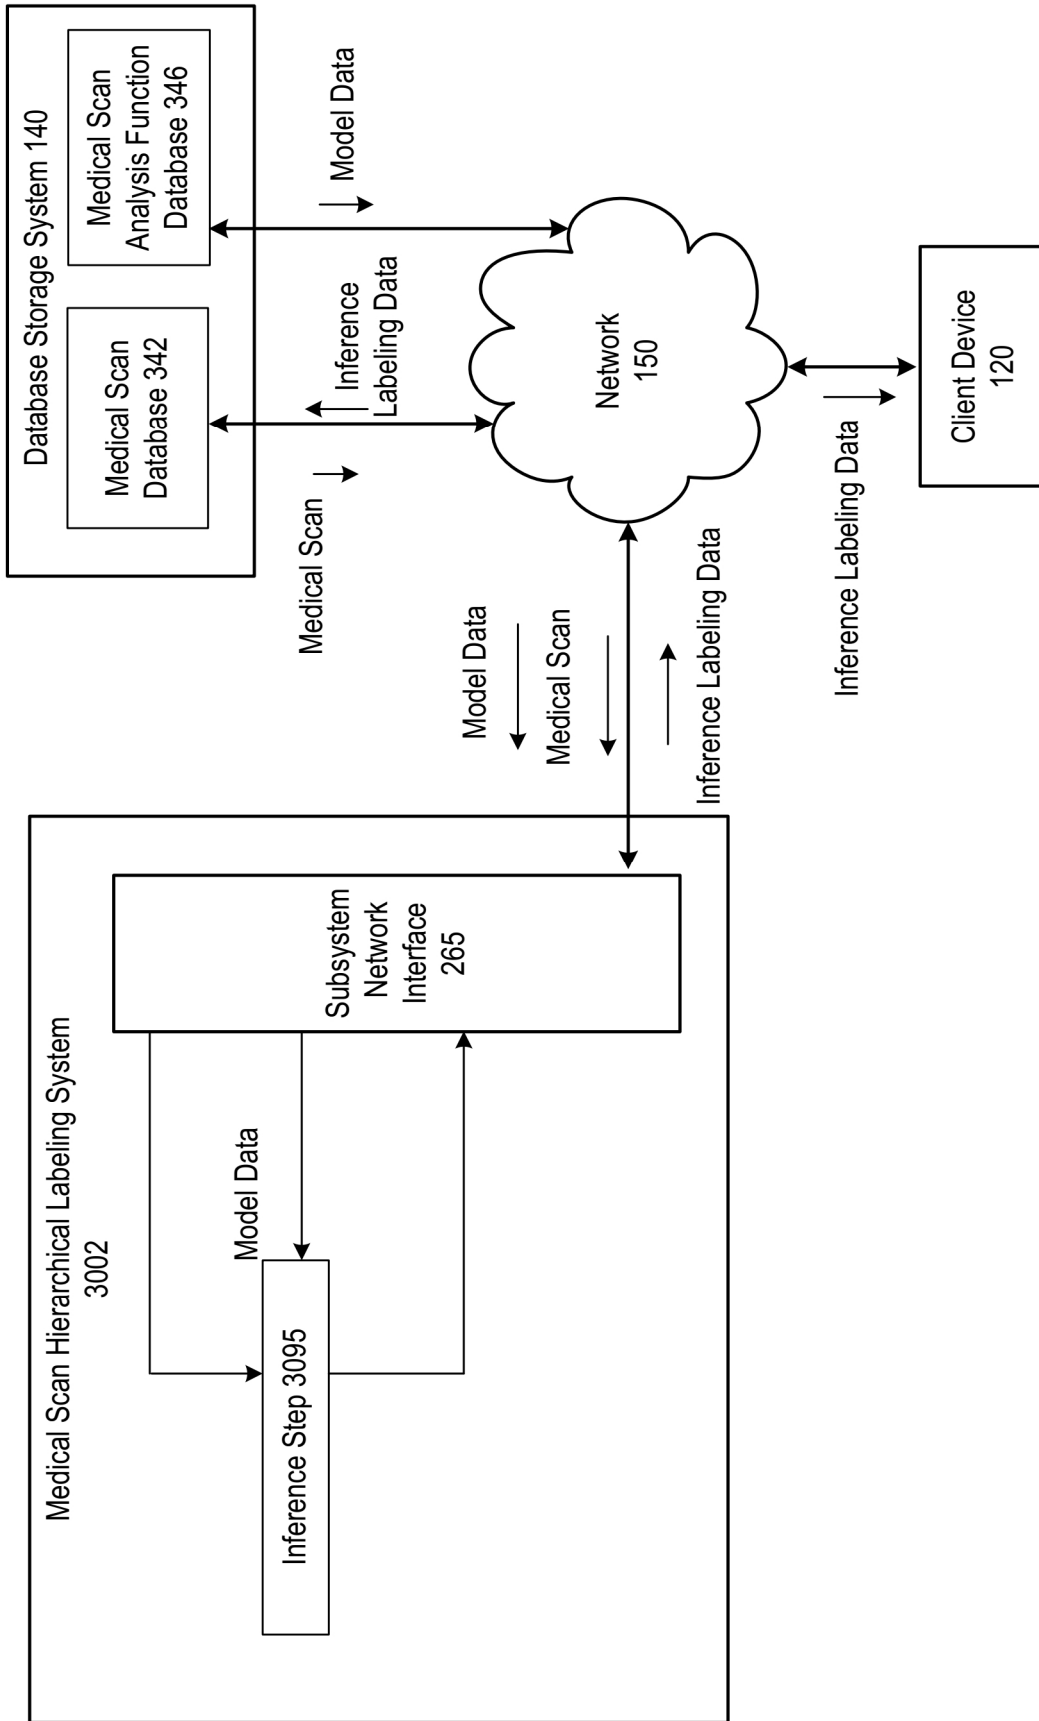


FIG. 12G

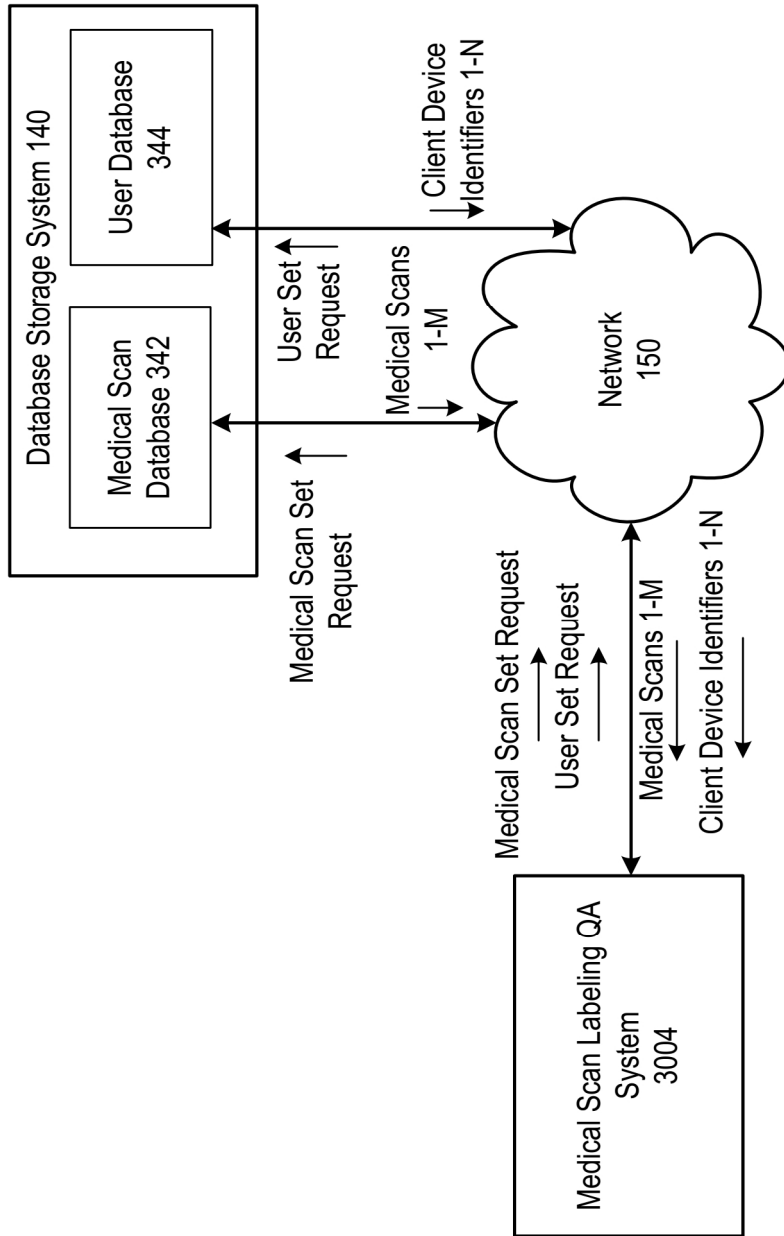


FIG. 13A

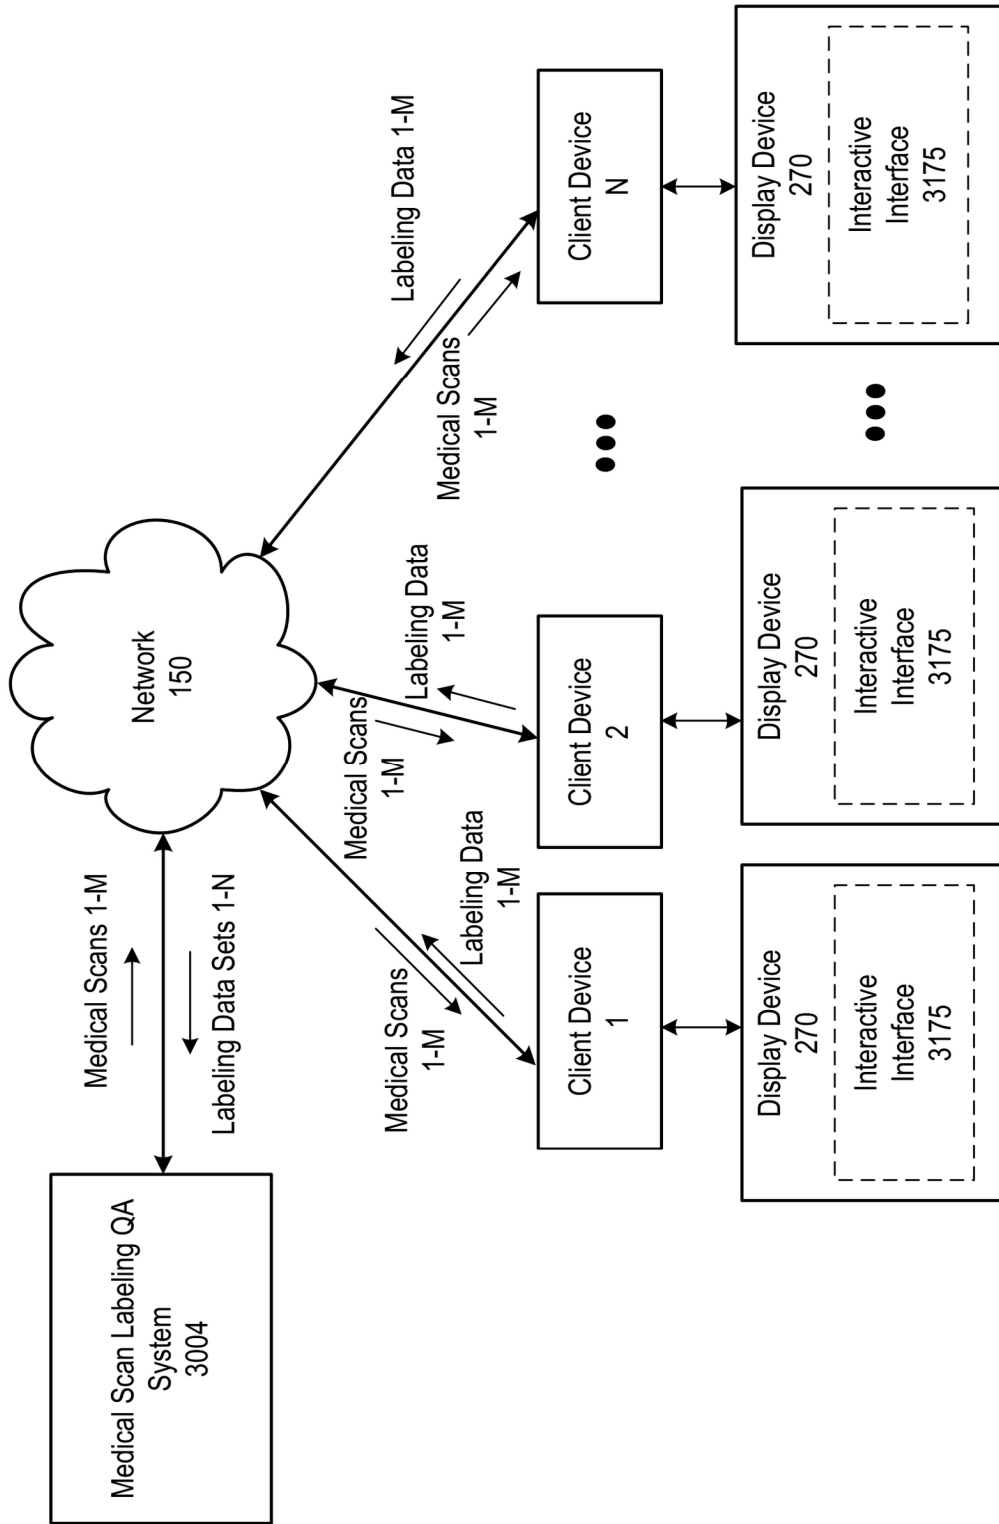


FIG. 13B

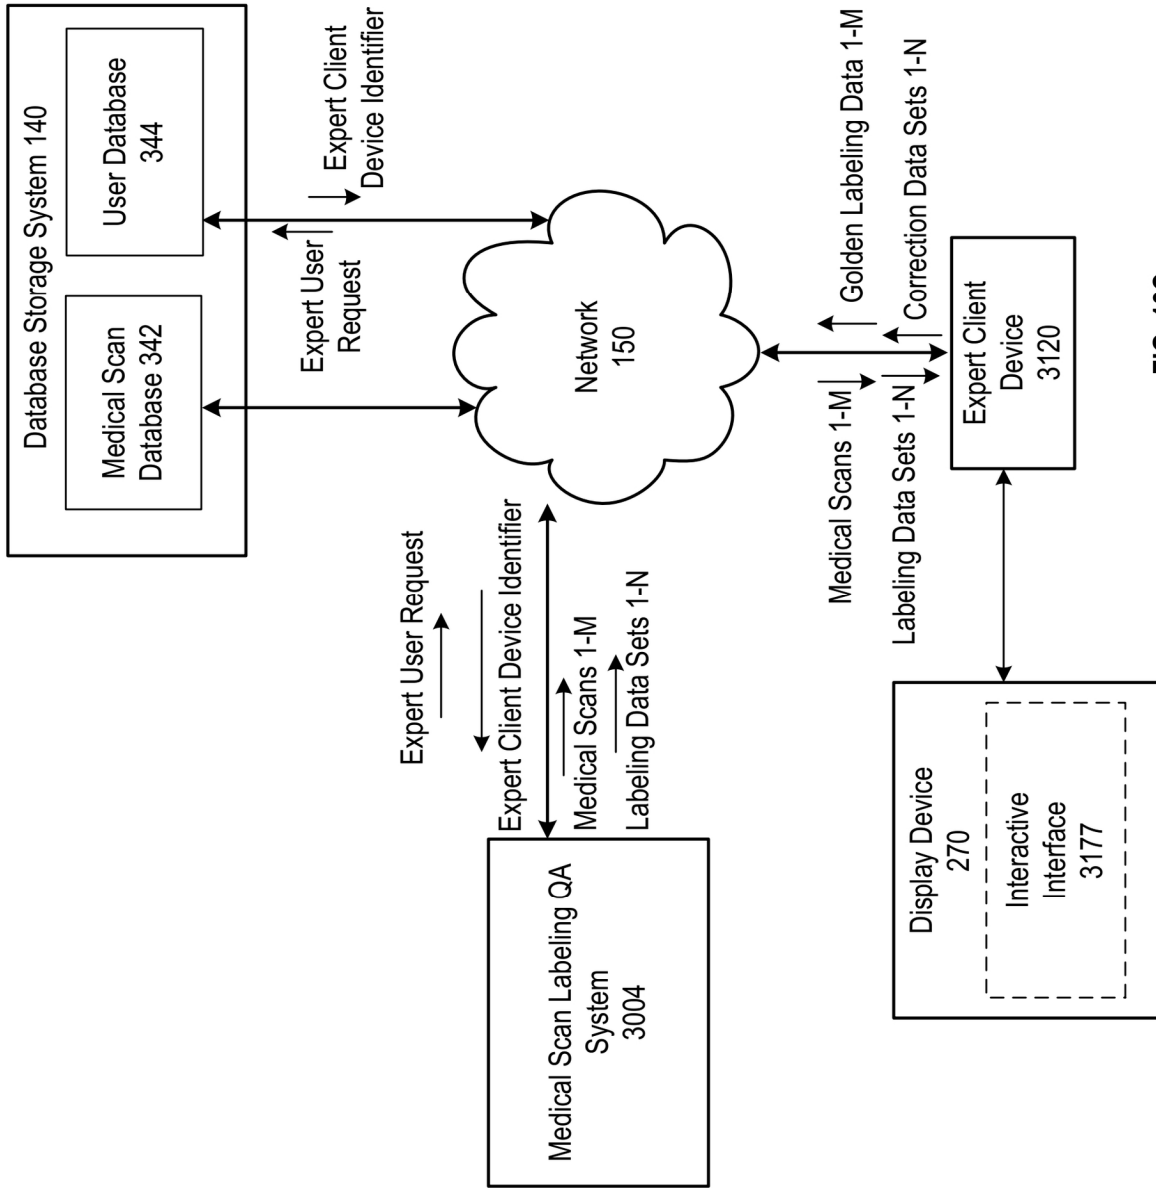


FIG. 13C

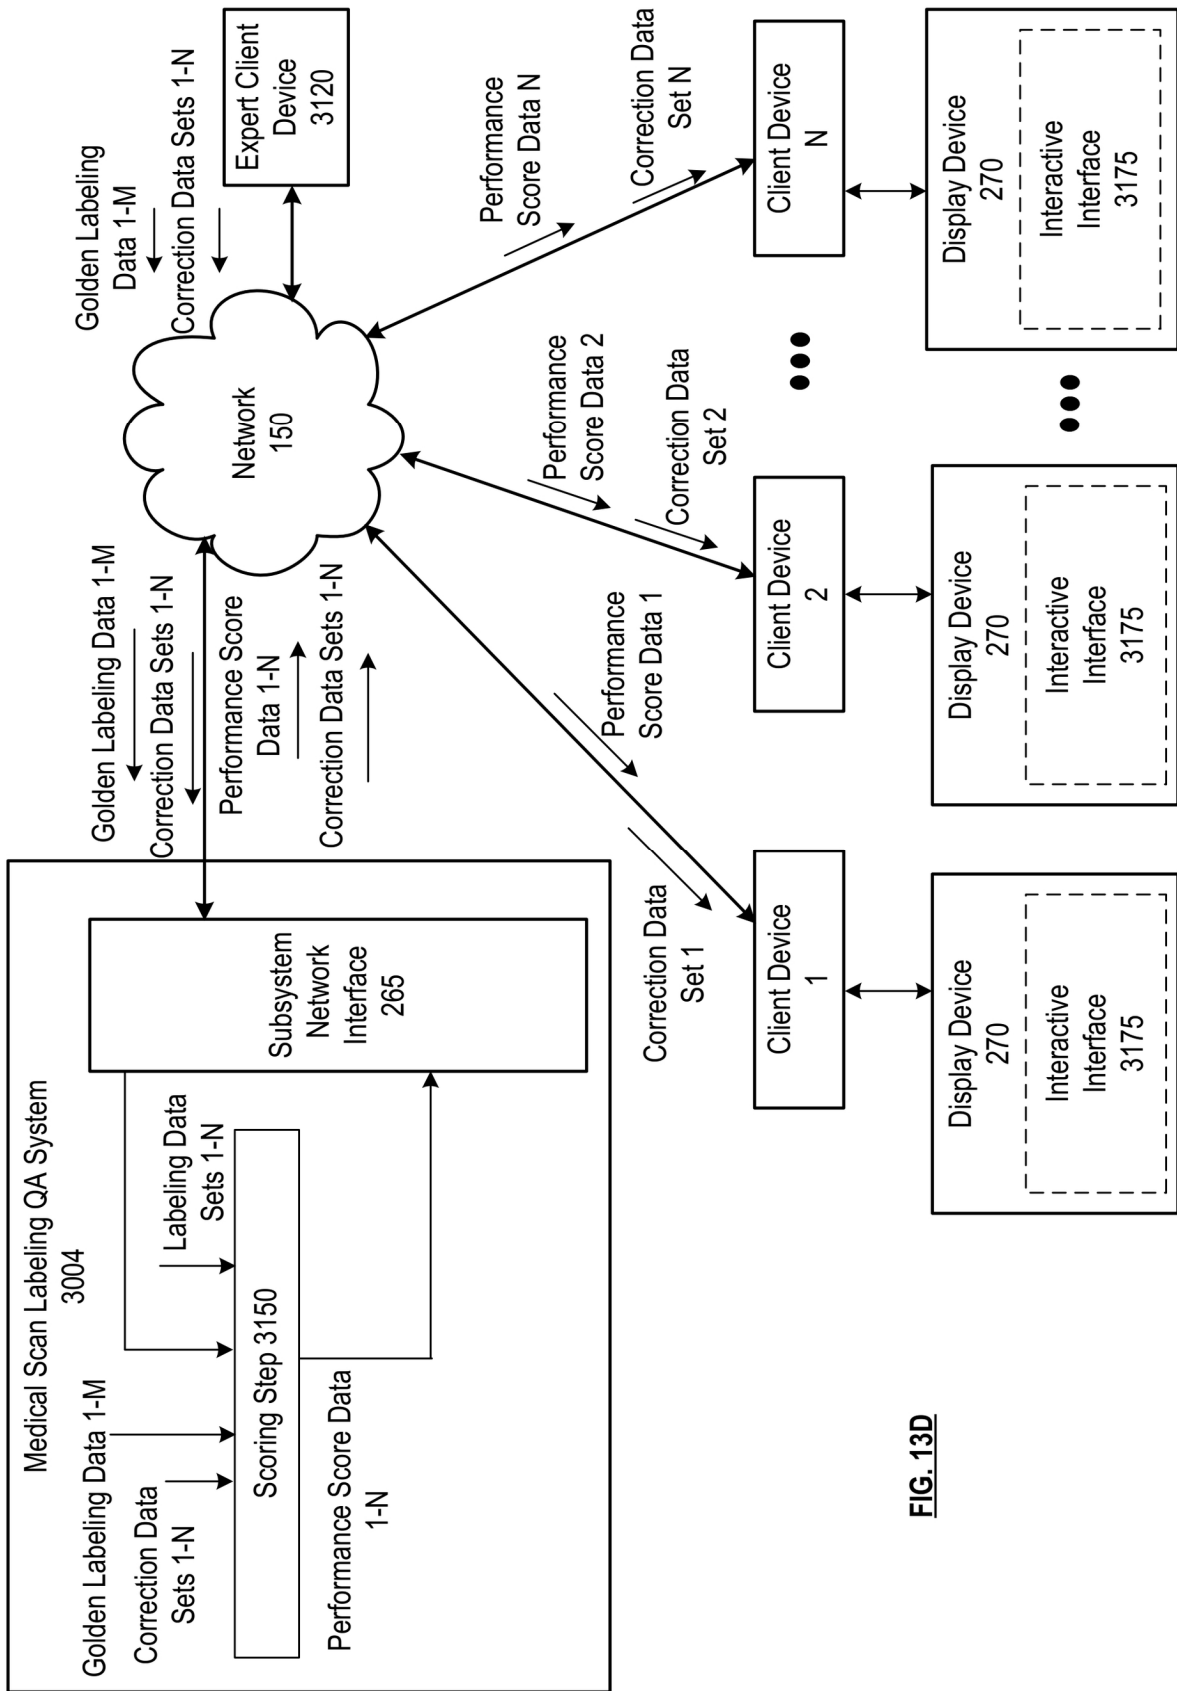


FIG. 13D




FIG. 13E

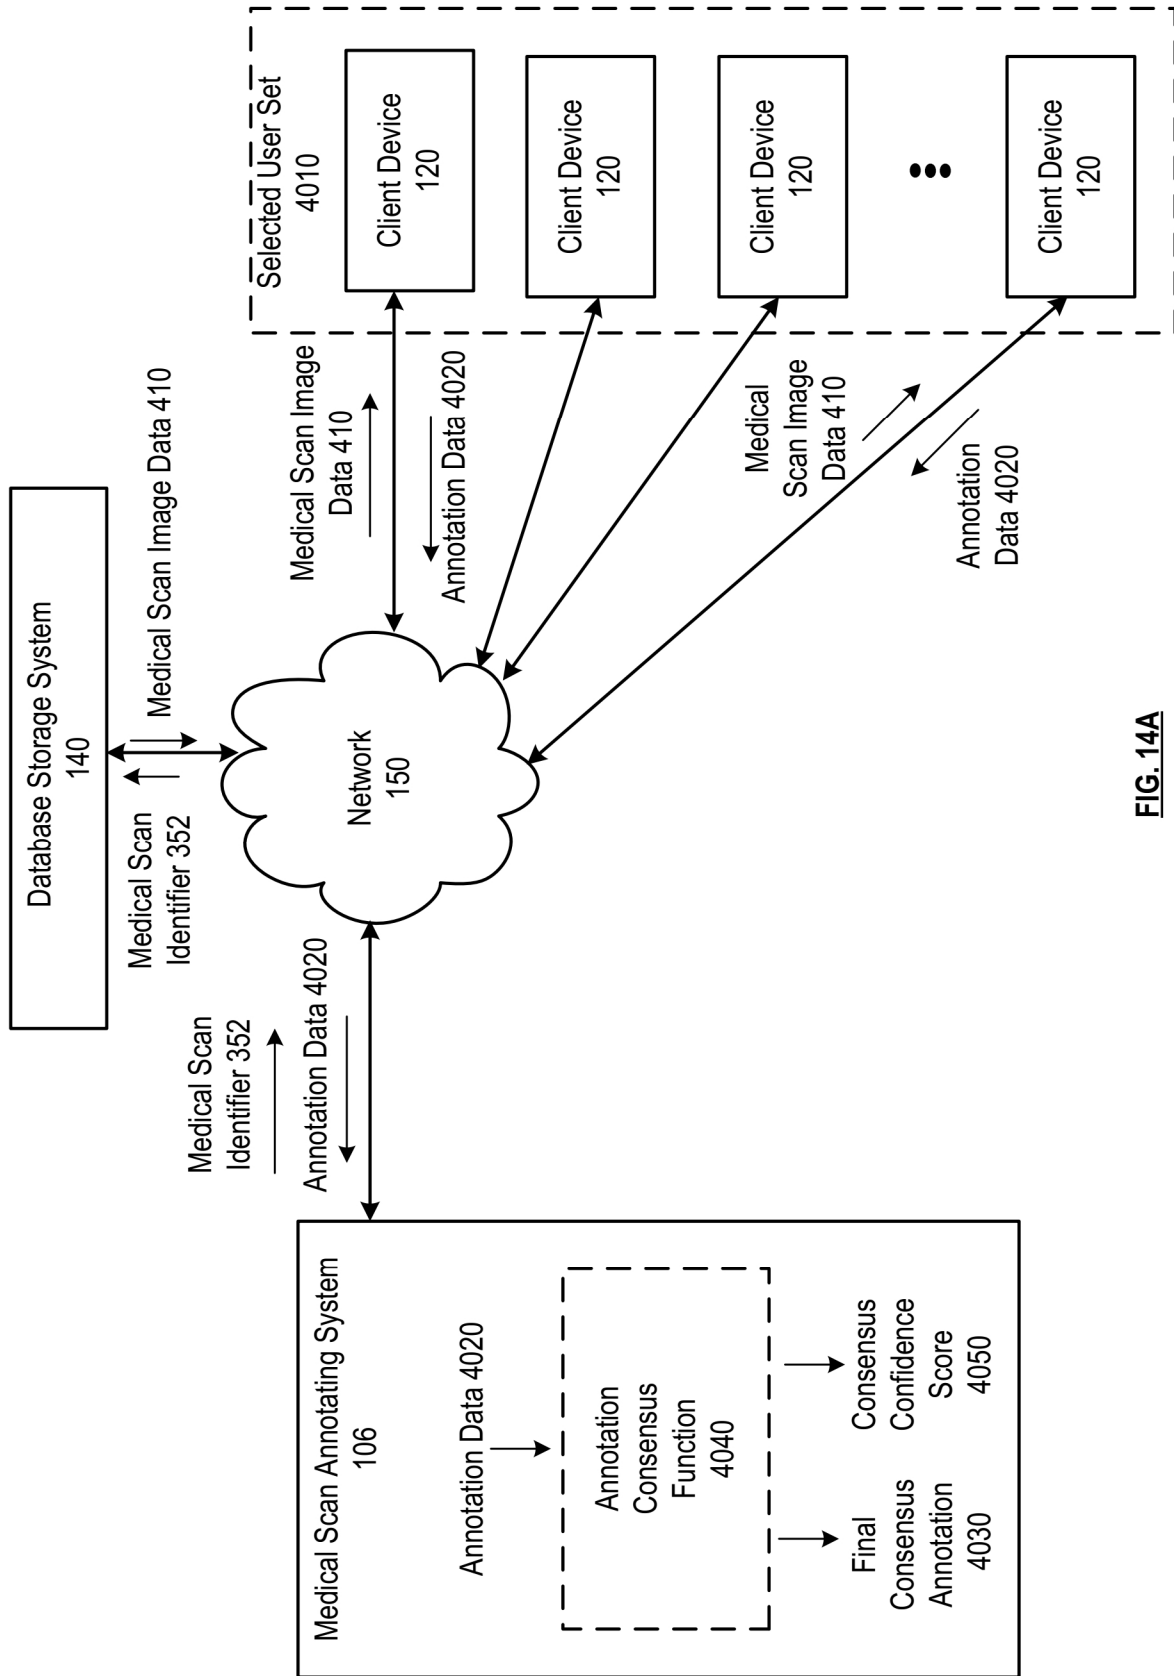


FIG. 14A




FIG. 14B

Submissions

1

Nodules

Please identify this nodule in a scan from each year in which it is visible, and indicate its ultimate malignancy.

Contours

Select

Select

Select

Did it become malignant?

No

Maybe

Yes

Please identify this nodule in a scan from each year in which it is visible, and indicate its ultimate malignancy.

Contours

Select

Select

T4: Invasive bronchiolo-alveolar adenocarcinoma of the left lower lobe (stage IA)
T6: Invasive squamous cell carcinoma, NOS of the left lower lobe (stage IA)

Please characterize each nodule displayed below

Nodule Characterizations

AnswerID of referenced nodule
33341

Does this nodule have spiculation?

Does this nodule have lobulation?

Nodule Texture

Solid

Nodule Calcification

absent

Nodule Sphericity

round

Nodule Internal Structure

soft

fluid

fat

air

T4: Invasive bronchiolo-alveolar adenocarcinoma of the left lower lobe (stage IA)
T6: Invasive squamous cell carcinoma, NOS of the left lower lobe (stage IA)

Please characterize each nodule displayed below

Nodule Characterizations

AnswerID of referenced nodule
33341

Does this nodule have spiculation?

Does this nodule have lobulation?

Nodule Texture

Solid

Part Solid/Mixed

Non-Solid/Ground Glass Opacity

Nodule Calcification

absent

Nodule Sphericity

round

Nodule Internal Structure

soft

fluid

fat

air

T4: Invasive bronchiolo-alveolar adenocarcinoma of the left lower lobe (stage IA)
T6: Invasive squamous cell carcinoma, NOS of the left lower lobe (stage IA)

Please characterize each nodule displayed below

Nodule Characterizations

AnswerID of referenced nodule
33341

Does this nodule have spiculation?

Does this nodule have lobulation?

Nodule Texture

popcorn

laminated

solid

Part Solid/Mixed

non-central

central

absent

stippled

Nodule Sphericity

round

Nodule Internal Structure

soft

fluid

fat

air

T4: Invasive bronchiolo-alveolar adenocarcinoma of the left lower lobe (stage IA)
T6: Invasive squamous cell carcinoma, NOS of the left lower lobe (stage IA)

Please characterize each nodule displayed below

Nodule Characterizations

AnswerID of referenced nodule
33341

Does this nodule have spiculation?

Does this nodule have lobulation?

Nodule Texture

Solid

Nodule Calcification

absent

Nodule Sphericity

linear

ovoid

round

Nodule Internal Structure

soft

fluid

fat

air

FIG. 14C

FIG. 14D

FIG. 14E

FIG. 14F

FIG. 14G

Nodules

Please identify this nodule in a scan from each year in which it is visible, and indicate its ultimate malignancy.

Contours

Select

Add new contour

Did it become malignant?

No Maybe Yes

Add new nodule

There's a problem

Please elaborate

Submit & Next

FIG. 14J

Nodules

Please identify this nodule in a scan from each year in which it is visible, and indicate its ultimate malignancy.

Contours

Add new contour

Did it become malignant?

No Maybe Yes

Add new nodule

There's a problem

Please elaborate

Submit & Next

FIG. 14I

Nodules

Add new nodule

There's a problem

Please elaborate

Submit & Next

FIG. 14H




FIG. 14K

In the context of a screening program on a generally healthy population, with no clinical history available, would you use a normal template for this study?

Please select an option

Submit & Next

FIG. 14L

In the context of a screening program on a generally healthy population, with no clinical history available, would you use a normal template for this study?

No, this study is abnormal.

Please add all abnormalities that you can identify in this study which you would include in your radiology report.

Please put an ROI the abnormality in all locations on all views where it is visible.

Add new ROI

Please select the category if this abnormality.

Please select an option

Add an abnormality

Submit & Next

FIG. 14N

In the context of a screening program on a generally healthy population, with no clinical history available, would you use a normal template for this study?

Please select an option

Yes, I would use a normal template for this study

No, this study is abnormal.

This is not a chest X-ray

FIG. 14M

Please select the category if this abnormality.

Please select an option

- Airway Abnormality (bronchiectasis or bronchiolitis)
- Anatomic Variant
- Atelectasis and/or scarring
- Consolidation
- Diaphragmatic Abnormality
- Enlarged Cardiac Contour
- Foreign Body (non-medical)
- Hyperlucent Thorax
- Mediastinal Abnormality
- Musculoskeletal Abnormality
- Nodule and/or mass
- Pleural Effusion
- Pleural Lesion
- Post Surgical Change
- Reticular Opacities
- Skin abnormality (mole, nipple shadow)
- Support devices and medical devices
- Abdominal abnormality not otherwise covered
- Other

FIG. 14O




FIG. 14P




FIG. 14Q

Assume patients A and B were part of a screening program on a healthy population, and no clinical history is available. If you had to give one patient a normal report, which would you choose?

- I would give Patient A a normal report.
- I would give Patient B a normal report.

Save Changes

Submissions

- 1
- 2
- 3
- 4

In the context of a screening program on a generally healthy population, with no clinical history available, would you use a normal template for this study?

No, this study is abnormal.

Please add all abnormalities that you can identify in this study which you would include in your radiology report.

Please put an ROI the abnormality in all locations on all views where it is visible.

Select

Please select the category if this abnormality.

Foreign Body (non-medical)

FIG. 14R

Please select all abnormalities

For each abnormality you identify in the study, please check the box below that best describes that abnormality.

Pulmonary Parenchyma
 Pulmonary Vasculature
 Plethora
 Diffuse
 Left
 Zone - Left Upper
 Zone - Left Mid
 Zone - Left Lower
 Lobe - Left Upper
 Lobe - Left Lower
 Right
 Focal/Multifocal
 Oligaemia
 Bronchial
 Abnormal Hilum
 Diaphragmatic Abnormality
 Pleural Abnormality
 Extra Pleural/Chest Wall
 Mediastinum
 Mass
 Compression of Structure
 Superior
 Inferior - Anterior
 Inferior - Middle
 Inferior - Posterior
 Heart
 Aorta
 Enlargement of Structure
 Ectasia
 Other
 Spine Abnormality
 Anatomic Variation
 Foreign Body
 Technical/Imaging Issue
 Other

Submit & Next

FIG. 14V

Please select all abnormalities

For each abnormality you identify in the study, please check the box below that best describes that abnormality.

Pulmonary Parenchyma
 Pulmonary Vasculature
 Plethora
 Diffuse
 Left
 Zone - Left Upper
 Zone - Left Mid
 Zone - Left Lower
 Lobe - Left Upper
 Lobe - Left Lower
 Right
 Focal/Multifocal
 Oligaemia
 Bronchial
 Abnormal Hilum
 Diaphragmatic Abnormality
 Pleural Abnormality
 Extra Pleural/Chest Wall
 Mediastinum
 Spine Abnormality
 Anatomic Variation
 Foreign Body
 Technical/Imaging Issue
 Other

Submit & Next

FIG. 14U

Please select all abnormalities

For each abnormality you identify in the study, please check the box below that best describes that abnormality.

Pulmonary Parenchyma
 Pulmonary Vasculature
 Plethora
 Diffuse
 Left
 Right
 Focal/Multifocal
 Oligaemia
 Bronchial
 Abnormal Hilum
 Diaphragmatic Abnormality
 Pleural Abnormality
 Extra Pleural/Chest Wall
 Mediastinum
 Spine Abnormality
 Anatomic Variation
 Foreign Body
 Technical/Imaging Issue
 Other

Submit & Next

FIG. 14T

Please select all abnormalities

For each abnormality you identify in the study, please check the box below that best describes that abnormality.

Pulmonary Parenchyma
 Pulmonary Vasculature
 Plethora
 Oligaemia
 Bronchial
 Abnormal Hilum
 Diaphragmatic Abnormality
 Pleural Abnormality
 Extra Pleural/Chest Wall
 Mediastinum
 Spine Abnormality
 Anatomic Variation
 Foreign Body
 Technical/Imaging Issue
 Other

Submit & Next

FIG. 14S

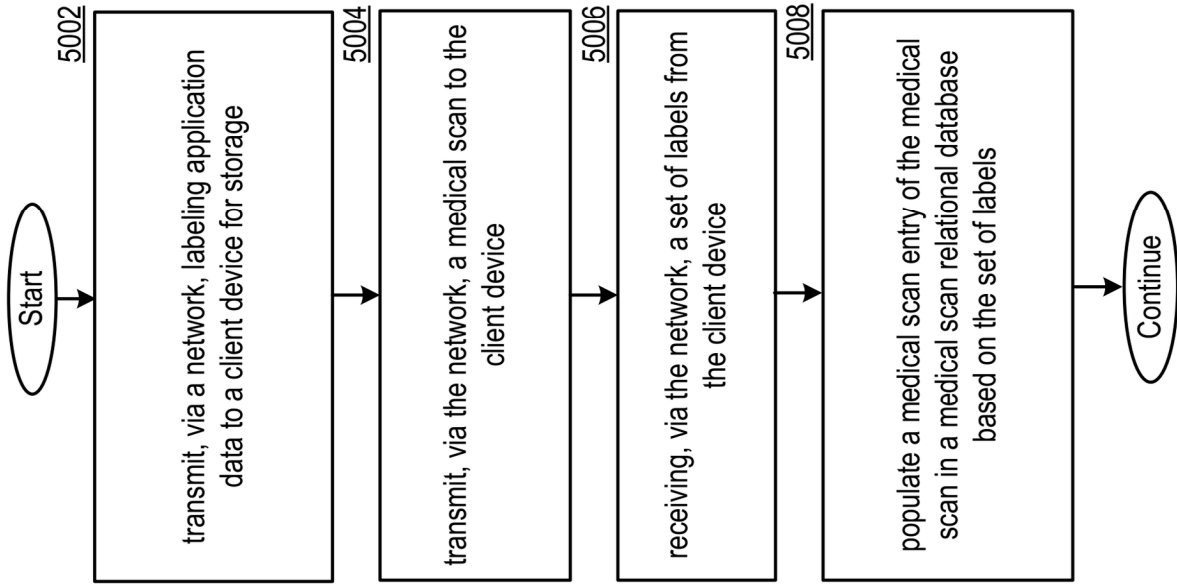


FIG. 15A

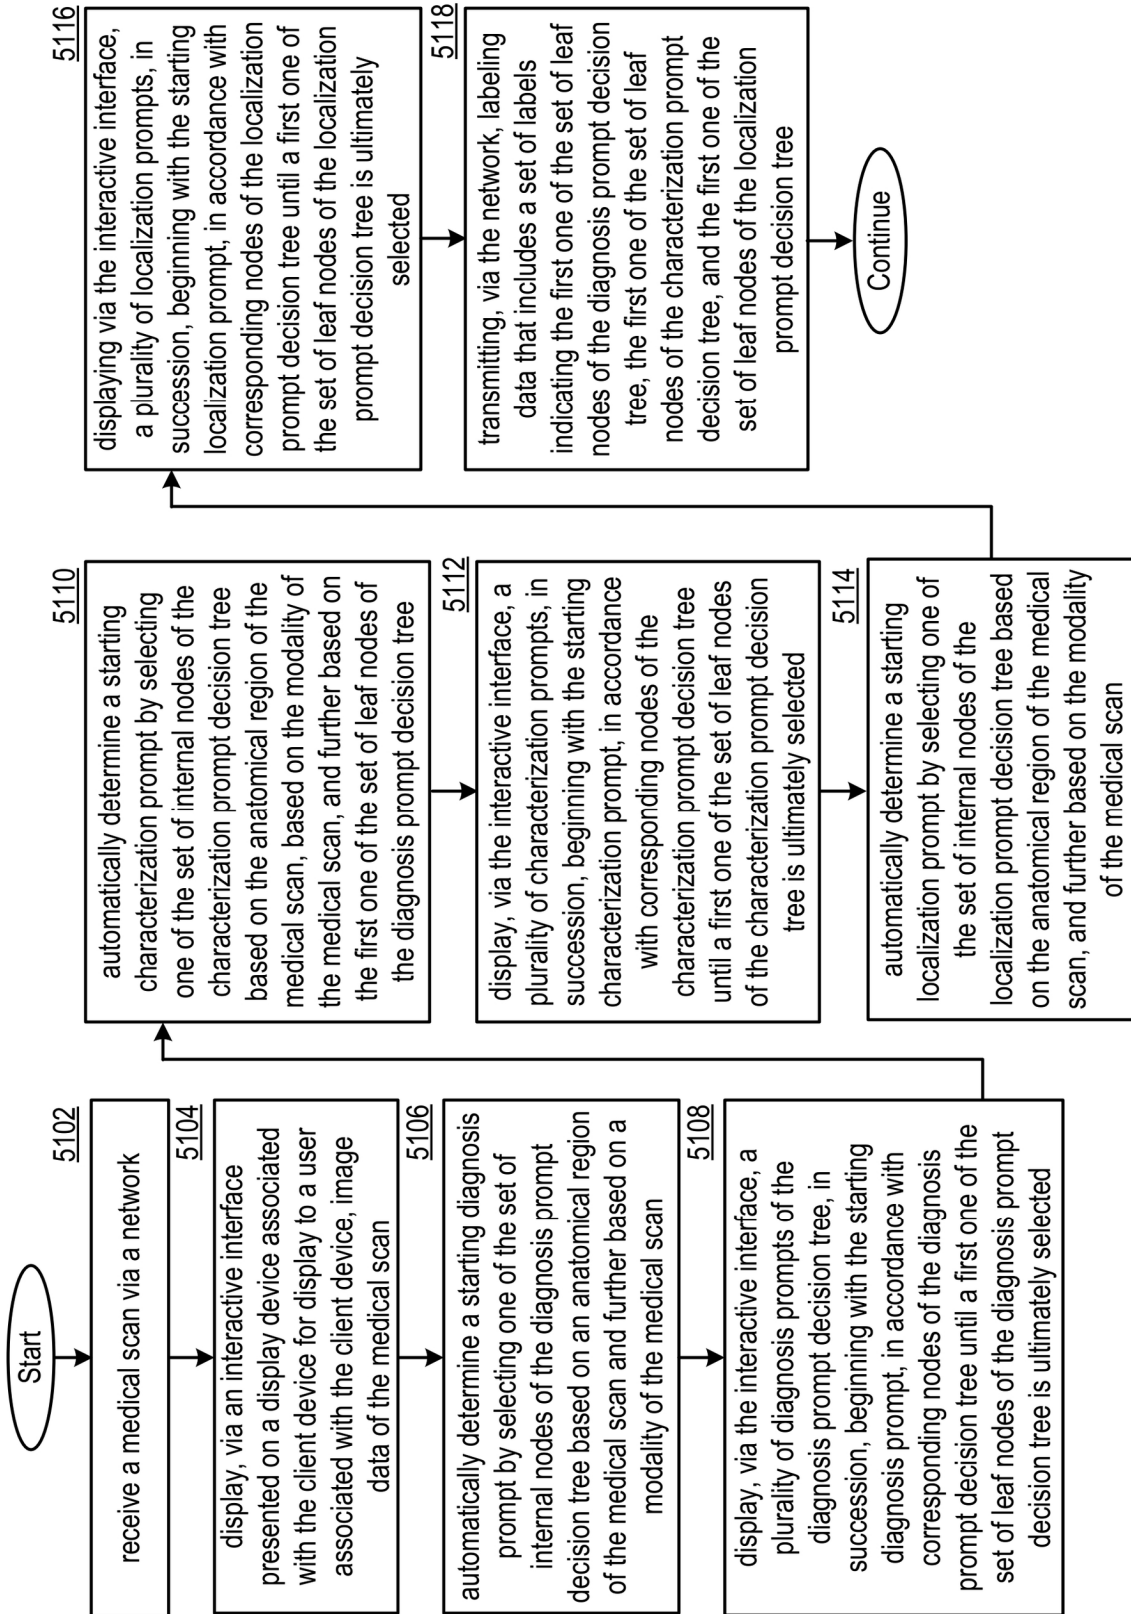


FIG. 15B

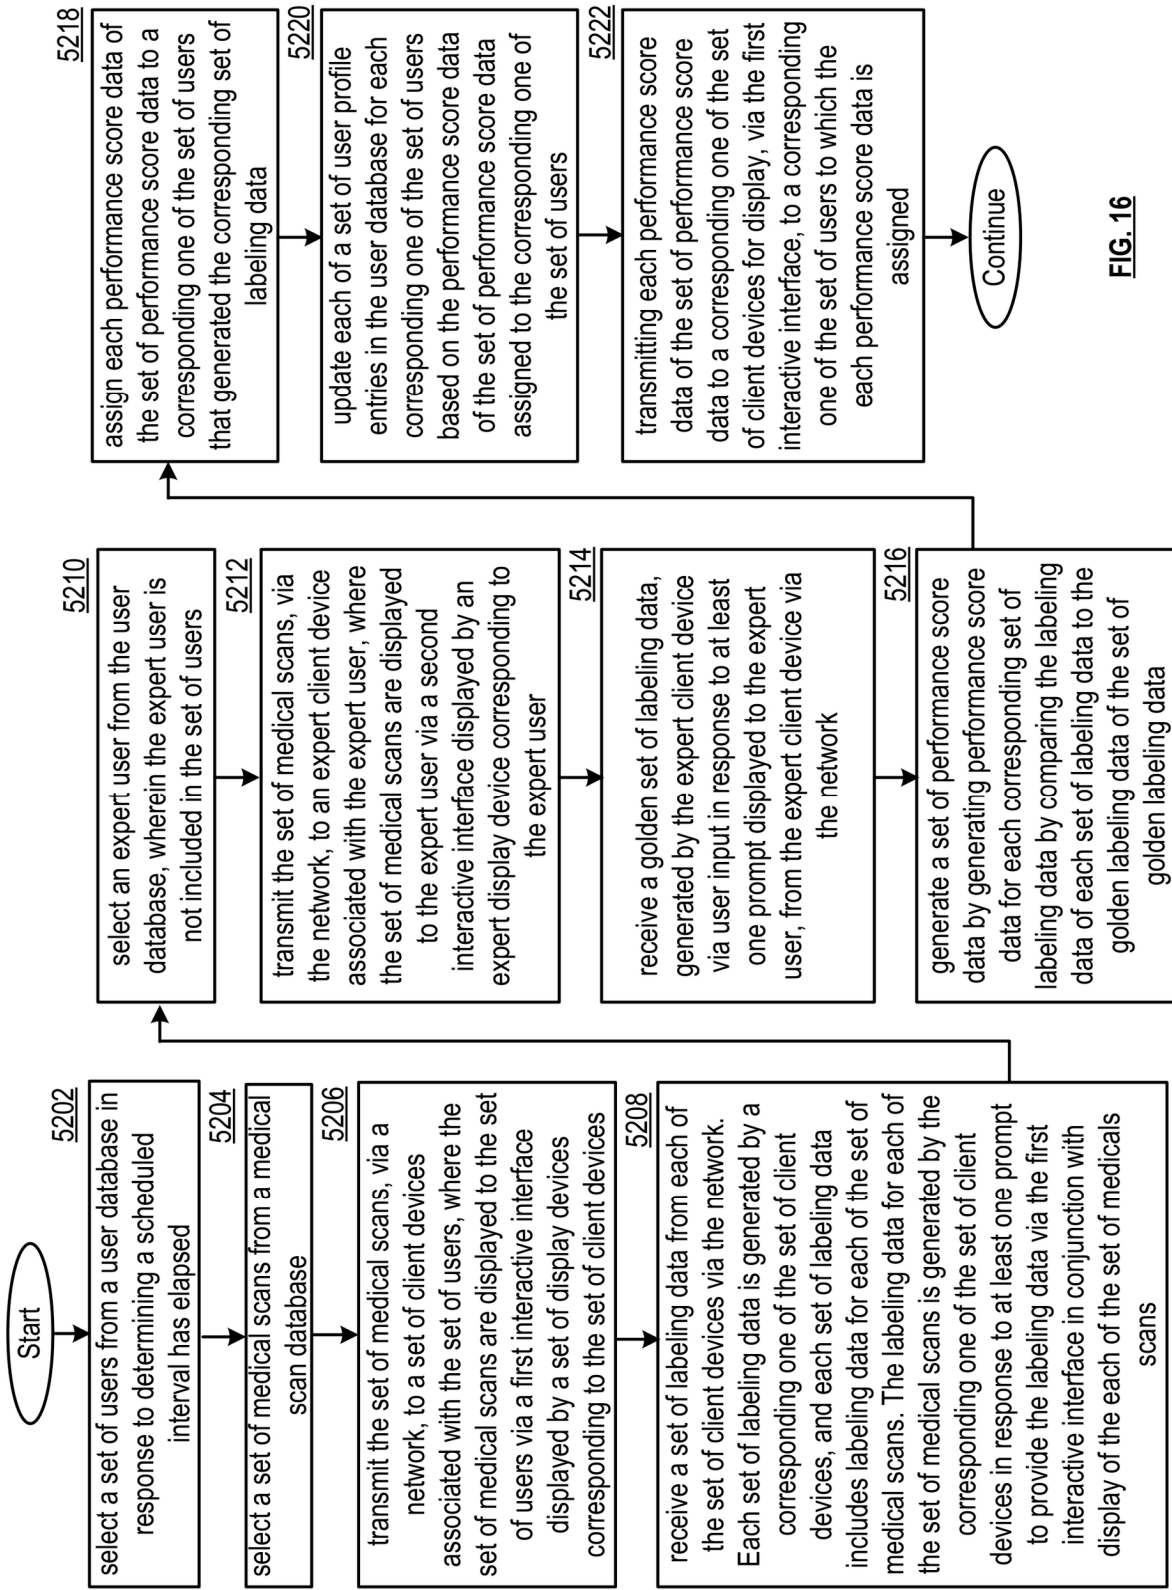


FIG. 16

1

MEDICAL SCAN LABELING QUALITY ASSURANCE SYSTEM AND METHODS FOR USE THEREWITH

CROSS REFERENCE TO RELATED APPLICATIONS

The present U.S. Utility Patent Application claims priority pursuant to 35 U.S.C. § 120 as a continuation of U.S. Utility application Ser. No. 16/363,019, entitled “MEDICAL SCAN LABELING QUALITY ASSURANCE SYSTEM”, filed Mar. 25, 2019, which claims priority pursuant to 35 U.S.C. § 119(e) to U.S. Provisional Application No. 62/770,334, entitled “LESION TRACKING SYSTEM”, filed Nov. 21, 2018, both of which are hereby incorporated herein by reference in their entirety and made part of the present U.S. Utility Patent Application for all purposes.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

Not applicable.

INCORPORATION-BY-REFERENCE OF MATERIAL SUBMITTED ON A COMPACT DISC

Not applicable.

BACKGROUND

Technical Field

This invention relates generally to medical imaging devices and knowledge-based systems used in conjunction with client/server network architectures.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING(S)

FIG. 1 is a schematic block diagram of an embodiment of a medical scan processing system;

FIG. 2A is a schematic block diagram of a client device in accordance with various embodiments;

FIG. 2B is a schematic block diagram of one or more subsystems in accordance with various embodiments;

FIG. 3 is a schematic block diagram of a database storage system in accordance with various embodiments;

FIG. 4A is schematic block diagram of a medical scan entry in accordance with various embodiments;

FIG. 4B is a schematic block diagram of abnormality data in accordance with various embodiments;

FIG. 5A is a schematic block diagram of a user profile entry in accordance with various embodiments;

FIG. 5B is a schematic block diagram of a medical scan analysis function entry in accordance with various embodiments;

FIGS. 6A-6B are schematic block diagram of a medical scan diagnosing system in accordance with various embodiments;

FIG. 7A is a flowchart representation of an inference step in accordance with various embodiments;

FIG. 7B is a flowchart representation of a detection step in accordance with various embodiments;

FIGS. 8A-8F are schematic block diagrams of a medical picture archive integration system in accordance with various embodiments;

2

FIG. 9 is a flowchart representation of a method for execution by a medical picture archive integration system in accordance with various embodiments;

FIG. 10A is a schematic block diagram of a de-identification system in accordance with various embodiments;

FIG. 10B is an illustration of an example of anonymizing patient identifiers in image data of a medical scan in accordance with various embodiments;

FIG. 11 presents a flowchart illustrating a method for execution by a de-identification system in accordance with various embodiments;

FIGS. 12A-12B are schematic block diagrams of a medical scan hierarchical labeling system in accordance with various embodiments;

FIG. 12C is an illustration of an example of a diagnosis prompt decision tree in accordance with various embodiments;

FIG. 12D is an illustration of an example of a characterization prompt decision tree in accordance with various embodiments;

FIG. 12E is an illustration of an example of a localization prompt decision tree in accordance with various embodiments;

FIGS. 12F-12G are schematic block diagrams of a medical scan hierarchical labeling system in accordance with various embodiments;

FIGS. 13A-13E are schematic block diagrams of a medical scan labeling quality assurance system in accordance with various embodiments;

FIGS. 14A-14B are schematic block diagrams of a medical scan annotating system in accordance with various embodiments;

FIGS. 14C-14V are graphical illustrations of an example interactive interface displayed on a client device in conjunction with various embodiments;

FIG. 15A presents a flowchart illustrating a method for execution by a medical scan hierarchical labeling system in accordance with various embodiments;

FIG. 15B presents a flowchart illustrating a method for execution by a client device in accordance with various embodiments; and

FIG. 16 presents a flowchart illustrating a method for execution by a medical scan labeling quality assurance system in accordance with various embodiments;

DETAILED DESCRIPTION

The present U.S. Utility Patent Application is related to U.S. Utility application Ser. No. 15/627,644, entitled “MEDICAL SCAN ASSISTED REVIEW SYSTEM”, filed 20 Jun. 2017, which claims priority pursuant to 35 U.S.C. § 119(e) to U.S. Provisional Application No. 62/511,150, entitled “MEDICAL SCAN ASSISTED REVIEW SYSTEM AND METHODS”, filed 25 May 2017, both of which are hereby incorporated herein by reference in their entirety and made part of the present U.S. Utility Patent Application for all purposes.

FIG. 1 presents a medical scan processing system **100**, which can include one or more medical scan subsystems **101** that communicate bidirectionally with one or more client devices **120** via a wired and/or wireless network **150**. The medical scan subsystems **101** can include a medical scan assisted review system **102**, medical scan report labeling system **104**, a medical scan annotator system **106**, a medical scan diagnosing system **108**, a medical scan interface feature evaluator system **110**, a medical scan image analysis system **112**, a medical scan natural language analysis system **114**,

and/or a medical scan comparison system **116**. Some or all of the subsystems **101** can utilize the same processing devices, memory devices, and/or network interfaces, for example, running on a same set of shared servers connected to network **150**. Alternatively or in addition, some or all of the subsystems **101** be assigned their own processing devices, memory devices, and/or network interfaces, for example, running separately on different sets of servers connected to network **150**. Some or all of the subsystems **101** can interact directly with each other, for example, where one subsystem's output is transmitted directly as input to another subsystem via network **150**. Network **150** can include one or more wireless and/or wired communication systems; one or more non-public intranet systems and/or public internet systems; and/or one or more local area networks (LAN) and/or wide area networks (WAN).

The medical scan processing system **100** can further include a database storage system **140**, which can include one or more servers, one or more memory devices of one or more subsystems **101**, and/or one or more other memory devices connected to network **150**. The database storage system **140** can store one or more shared databases and/or one or more files stored on one or more memory devices that include database entries as described herein. The shared databases and/or files can each be utilized by some or all of the subsystems of the medical scan processing system, allowing some or all of the subsystems and/or client devices to retrieve, edit, add, or delete entries to the one or more databases and/or files.

The one or more client devices **120** can each be associated with one or more users of one or more subsystems of the medical scan processing system. Some or all of the client devices can be associated with hospitals or other medical institutions and/or associated with medical professionals, employees, or other individual users for example, located at one or more of the medical institutions. Some of the client devices **120** can correspond to one or more administrators of one or more subsystems of the medical scan processing system, allowing administrators to manage, supervise, or override functions of one or more subsystems for which they are responsible.

Some or all of the subsystems **101** of the medical scan processing system **100** can include a server that presents a website for operation via a browser of client devices **120**. Alternatively or in addition, each client device can store application data corresponding to some or all subsystems, for example, a subset of the subsystems that are relevant to the user in a memory of the client device, and a processor of the client device can display the interactive interface based on instructions in the interface data stored in memory. For example, the website presented by a subsystem can operate via the application. Some or all of the websites presented can correspond to multiple subsystems, for example, where the multiple subsystems share the server presenting the website. Furthermore, the network **150** can be configured for secure and/or authenticated communications between the medical scan subsystems **101**, the client devices **120** and the database storage system **140** to protect the data stored in the database storage system and the data communicated between the medical scan subsystems **101**, the client devices **120** and the database storage system **140** from unauthorized access.

The medical scan assisted review system **102** can be used to aid medical professionals or other users in diagnosing, triaging, classifying, ranking, and/or otherwise reviewing medical scans by presenting a medical scan for review by a user by transmitting medical scan data of a selected medical scan and/or interface feature data of selected interface

features of to a client device **120** corresponding to a user of the medical scan assisted review system for display via a display device of the client device. The medical scan assisted review system **102** can generate scan review data for a medical scan based on user input to the interactive interface displayed by the display device in response to prompts to provide the scan review data, for example, where the prompts correspond to one or more interface features.

The medical scan assisted review system **102** can be operable to receive, via a network, a medical scan for review. Abnormality annotation data can be generated by identifying one or more of abnormalities in the medical scan by utilizing a computer vision model that is trained on a plurality of training medical scans. The abnormality annotation data can include location data and classification data for each of the plurality of abnormalities and/or data that facilitates the visualization of the abnormalities in the scan image data. Report data including text describing each of the plurality of abnormalities is generated based on the abnormality data. The visualization and the report data, which can collectively be displayed annotation data, can be transmitted to a client device. A display device associated with the client device can display the visualization in conjunction with the medical scan via an interactive interface, and the display device can further display the report data via the interactive interface.

In various embodiments, longitudinal data, such as one or more additional scans of longitudinal data **433** of the medical scan or of similar scans, can be displayed in conjunction with the medical scan automatically, or in response to the user electing to view longitudinal data via user input. For example, the medical scan assisted review system can retrieve a previous scan or a future scan for the patient from a patient database or from the medical scan database automatically or in response to the user electing to view past patient data. One or more previous scans can be displayed in one or more corresponding windows adjacent to the current medical scan. For example, the user can select a past scan from the longitudinal data for display. Alternatively or in addition, the user can elect longitudinal parameters such as amount of time elapsed, scan type, electing to select the most recent and/or least recent scan, electing to select a future scan, electing to select a scan at a date closest to the scan, or other criteria, and the medical scan assisted review system can automatically select a previous scan that compares most favorably to the longitudinal parameters. The selected additional scan can be displayed in an adjacent window alongside the current medical scan. In some embodiments, multiple additional scans will be selected and can be displayed in multiple adjacent windows.

In various embodiments, a first window displaying an image slice **412** of the medical scan and an adjacent second window displaying an image slice of a selected additional scan will display image slices **412** determined to correspond with the currently displayed slice **412** of the medical scan. As described with respect to selecting a slice of a selected similar medical scan for display, this can be achieved based on selecting the image slice with a matching slice number, based on automatically determining the image slice that most closely matches the anatomical region corresponding to the currently displayed slice of the current scan, and/or based on determining the slice in the previous scan with the most similar view of the abnormality as the currently displayed slice. The user can use a single scroll bar or other single user input indication to jump to a different image slice, and the multiple windows can simultaneously display the same numbered image slice, or can scroll or jump by the same number of slices if different slice numbers are initially

displayed. In some embodiments, three or more adjacent windows corresponding to the medical scan and two or more additional scans are displayed, and can all be controlled with the single scroll bar in a similar fashion.

The medical scan assisted review system **102** can automatically detect previous states of the identified abnormalities based on the abnormality data, such as the abnormality location data. The detected previous states of the identified abnormality can be circled, highlighted, or otherwise indicated in their corresponding window. The medical scan assisted review system **102** can retrieve classification data for the previous state of the abnormality by retrieving abnormality annotation data **442** of the similar abnormality mapped to the previous scan from the medical scan database **342**. This data may not be assigned to the previous scan, and the medical scan assisted review system can automatically determine classification or other diagnosis data for the previous medical scan by utilizing the medical scan image analysis system as discussed. Alternatively or in addition, some or all of the abnormality classification data **445** or other diagnosis data **440** for the previous scan can be assigned values determined based on the abnormality classification data or other diagnosis data determined for the current scan. Such abnormality classification data **445** or other diagnosis data **440** determined for the previous scan can be mapped to the previous scan, and or mapped to the longitudinal data **433**, in the database and/or transmitted to a responsible entity via the network.

The medical assisted review system can automatically generate state change data such as a change in size, volume, malignancy, or other changes to various classifiers of the abnormality. This can be achieved by automatically comparing image data of one or more previous scans and the current scan and/or by comparing abnormality data of the previous scan to abnormality data of the current scan. In some embodiments, such metrics can be calculated by utilizing the medical scan similarity analysis function, for example, where the output of the medical scan similarity analysis function such as the similarity score indicates distance, error, or other measured discrepancy in one or more abnormality classifier categories **444** and/or abnormality pattern categories **446**. This calculated distance, error, or other measured discrepancy in each category can be used to quantify state change data, indicate a new classifier in one or more categories, to determine if a certain category has become more or less severe, or otherwise determine how the abnormality has changed over time. In various embodiments, this data can be displayed in one window, for example, where an increase in abnormality size is indicated by overlaying or highlighting an outline of the current abnormality over the corresponding image slice of the previous abnormality, or vice versa. In various embodiments where several past scans are available, such state change data can be determined over time, and statistical data showing growth rate changes over time or malignancy changes over time can be generated, for example, indicating if a growth rate is lessening or worsening over time. Image slices corresponding to multiple past scans can be displayed in sequence, for example, where a first scroll bar allows a user to scroll between image slice numbers, and a second scroll bar allows a user to scroll between the same image slice over time. In various embodiments the abnormality data, heat map data, or other interface features will be displayed in conjunction with the image slices of the past image data.

The medical scan report labeling system **104** can be used to automatically assign medical codes to medical scans

based on user identified keywords, phrases, or other relevant medical condition terms of natural text data in a medical scan report of the medical scan, identified by users of the medical scan report labeling system **104**. The medical scan report labeling system **104** can be operable to transmit a medical report that includes natural language text to a first client device for display. Identified medical condition term data can be received from the first client device in response. An alias mapping pair in a medical label alias database can be identified by determining that a medical condition term of the alias mapping pair compares favorably to the identified medical condition term data. A medical code that corresponds to the alias mapping pair and a medical scan that corresponds to the medical report can be transmitted to a second client device of an expert user for display, and accuracy data can be received from the second client device in response. The medical code is mapped to the first medical scan in a medical scan database when the accuracy data indicates that the medical code compares favorably to the medical scan.

The medical scan annotator system **106** can be used to gather annotations of medical scans based on review of the medical scan image data by users of the system such as radiologists or other medical professionals. Medical scans that require annotation, for example, that have been triaged from a hospital or other triaging entity, can be sent to multiple users selected by the medical scan annotator system **106**, and the annotations received from the multiple medical professionals can be processed automatically by a processing system of the medical scan annotator system, allowing the medical scan annotator system to automatically determine a consensus annotation of each medical scan. Furthermore, the users can be automatically scored by the medical scan annotator system based on how closely their annotation matches to the consensus annotation or some other truth annotation, for example, corresponding to annotations of the medical scan assigned a truth flag. Users can be assigned automatically to annotate subsequent incoming medical scans based on their overall scores and/or based on categorized scores that correspond to an identified category of the incoming medical scan.

The medical scan annotator system **106** can be operable to select a medical scan for transmission via a network to a first client device and a second client device for display via an interactive interface, and annotation data can be received from the first client device and the second client device in response. Annotation similarity data can be generated by comparing the first annotation data to the second annotation data, and consensus annotation data can be generated based on the first annotation data and the second annotation data in response to the annotation similarity data indicating that the difference between the first annotation data and the second annotation data compares favorably to an annotation discrepancy threshold. The consensus annotation data can be mapped to the medical scan in a medical scan database.

A medical scan diagnosing system **108** can be used by hospitals, medical professionals, or other medical entities to automatically produce inference data for given medical scans by utilizing computer vision techniques and/or natural language processing techniques. This automatically generated inference data can be used to generate and/or update diagnosis data or other corresponding data of corresponding medical scan entries in a medical scan database. The medical scan diagnosing system can utilize a medical scan database, user database, and/or a medical scan analysis function database by communicating with the database storage sys-

tem **140** via the network **150**, and/or can utilize another medical scan database, user database, and/or function database stored in local memory.

The medical scan diagnosing system **108** can be operable to receive a medical scan. Diagnosis data of the medical scan can be generated by performing a medical scan inference function on the medical scan. The first medical scan can be transmitted to a first client device associated with a user of the medical scan diagnosing system in response to the diagnosis data indicating that the medical scan corresponds to a non-normal diagnosis. The medical scan can be displayed to the user via an interactive interface displayed by a display device corresponding to the first client device. Review data can be received from the first client device, where the review data is generated by the first client device in response to a prompt via the interactive interface. Updated diagnosis data can be generated based on the review data. The updated diagnosis data can be transmitted to a second client device associated with a requesting entity.

A medical scan interface feature evaluating system **110** can be used evaluate proposed interface features or currently used interface features of an interactive interface to present medical scans for review by medical professionals or other users of one or more subsystems **101**. The medical scan interface feature evaluator system **110** can be operable to generate an ordered image-to-prompt mapping by selecting a set of user interface features to be displayed with each of an ordered set of medical scans. The set of medical scans and the ordered image-to-prompt mapping can be transmitted to a set of client devices. A set of responses can be generated by each client device in response to sequentially displaying each of the set of medical scans in conjunction with a mapped user interface feature indicated in the ordered image-to-prompt mapping via a user interface. Response score data can be generated by comparing each response to truth annotation data of the corresponding medical scan. Interface feature score data corresponding to each user interface feature can be generated based on aggregating the response score data, and is used to generate a ranking of the set of user interface features.

A medical scan image analysis system **112** can be used to generate and/or perform one or more medical scan image analysis functions by utilizing a computer vision-based learning algorithm **1350** on a training set of medical scans with known annotation data, diagnosis data, labeling and/or medical code data, report data, patient history data, patient risk factor data, and/or other metadata associated with medical scans. These medical scan image analysis functions can be used to generate inference data for new medical scans that are triaged or otherwise require inferred annotation data, diagnosis data, labeling and/or medical code data, and/or report data. For example, some medical scan image analysis functions can correspond to medical scan inference functions of the medical scan diagnosing system or other medical scan analysis functions of a medical scan analysis function database. The medical scan image analysis functions can be used to determine whether or not a medical scan is normal, to detect the location of an abnormality in one or more slices of a medical scan, and/or to characterize a detected abnormality. The medical scan image analysis system can be used to generate and/or perform computer vision based medical scan image analysis functions utilized by other subsystems of the medical scan processing system as described herein, aiding medical professionals to diagnose patients and/or to generate further data and models to characterize medical scans. The medical scan image analysis system can include a processing system that includes a processor and a memory

that stores executable instructions that, when executed by the processing system, facilitate performance of operations.

The medical scan image analysis system **112** can be operable to receive a plurality of medical scans that represent a three-dimensional anatomical region and include a plurality of cross-sectional image slices. A plurality of three-dimensional subregions corresponding to each of the plurality of medical scans can be generated by selecting a proper subset of the plurality of cross-sectional image slices from each medical scan, and by further selecting a two-dimensional subregion from each proper subset of cross-sectional image slices. A learning algorithm can be performed on the plurality of three-dimensional subregions to generate a neural network. Inference data corresponding to a new medical scan received via the network can be generated by performing an inference algorithm on the new medical scan by utilizing the neural network. An inferred abnormality can be identified in the new medical scan based on the inference data.

The medical scan natural language analysis system **114** can determine a training set of medical scans with medical codes determined to be truth data. Corresponding medical reports and/or other natural language text data associated with a medical scan can be utilized to train a medical scan natural language analysis function by generating a medical report natural language model. The medical scan natural language analysis function can be utilized to generate inference data for incoming medical reports for other medical scans to automatically determine corresponding medical codes, which can be mapped to corresponding medical scans. Medical codes assigned to medical scans by utilizing the medical report natural language model can be utilized by other subsystems, for example, to train other medical scan analysis functions, to be used as truth data to verify annotations provided via other subsystems, to aid in diagnosis, or otherwise be used by other subsystems as described herein.

A medical scan comparison system **116** can be utilized by one or more subsystems to identify and/or display similar medical scans, for example, to perform or determine function parameters for a medical scan similarity analysis function, to generate or retrieve similar scan data, or otherwise compare medical scan data. The medical scan comparison system **116** can also utilize some or all features of other subsystems as described herein. The medical scan comparison system **116** can be operable to receive a medical scan via a network and can generate similar scan data. The similar scan data can include a subset of medical scans from a medical scan database and can be generated by performing an abnormality similarity function, such as medical scan similarity analysis function, to determine that a set of abnormalities included in the subset of medical scans compare favorably to an abnormality identified in the medical scan. At least one cross-sectional image can be selected from each medical scan of the subset of medical scans for display on a display device associated with a user of the medical scan comparison system in conjunction with the medical scan.

FIG. 2A presents an embodiment of client device **120**. Each client device **120** can include one or more client processing devices **230**, one or more client memory devices **240**, one or more client input devices **250**, one or more client network interfaces **260** operable to more support one or more communication links via the network **150** indirectly and/or directly, and/or one or more client display devices **270**, connected via bus **280**. Client applications **202**, **204**, **206**, **208**, **210**, **212**, **214**, and/or **216** correspond to subsystems **102**, **104**, **106**, **108**, **110**, **112**, **114**, and/or **116** of the

medical scan processing system respectfully. Each client device **120** can receive the application data from the corresponding subsystem via network **150** by utilizing network interface **260**, for storage in the one or more memory devices **240**. In various embodiments, some or all client devices **120** can include a computing device associated with a radiologist, medical entity, or other user of one or more subsystems as described herein.

The one or more processing devices **230** can display interactive interface **275** on the one or more client display devices **270** in accordance with one or more of the client applications **202, 204, 206, 208, 210, 212, 214,** and/or **216**, for example, where a different interactive interface **275** is displayed for some or all of the client applications in accordance with the website presented by the corresponding subsystem **102, 104, 106, 108, 110, 112, 114** and/or **116**. The user can provide input in response to menu data or other prompts presented by the interactive interface via the one or more client input devices **250**, which can include a microphone, mouse, keyboard, touchscreen or display device **270** itself or other touchscreen, and/or other device allowing the user to interact with the interactive interface. The one or more processing devices **230** can process the input data and/or send raw or processed input data to the corresponding subsystem, and/or can receive and/or generate new data in response for presentation via the interactive interface **275** accordingly, by utilizing network interface **260** to communicate bidirectionally with one or more subsystems and/or databases of the medical scan processing system via network **150**.

FIG. **2B** presents an embodiment of a subsystem **101**, which can be utilized in conjunction with subsystem **102, 104, 106, 108, 110, 112, 114** and/or **116**. Each subsystem **101** can include one or more subsystem processing devices **235**, one or more subsystem memory devices **245**, and/or one or more subsystem network interfaces **265**, connected via bus **285**. The subsystem memory devices **245** can store executable instructions that, when executed by the one or more subsystem processing devices **235**, facilitate performance of operations by the subsystem **101**, as described for each subsystem herein.

FIG. **3** presents an embodiment of the database storage system **140**. Database storage system **140** can include at least one database processing device **330**, at least one database memory device **340**, and at least one database network interface **360**, operable to more support one or more communication links via the network **150** indirectly and/or directly, all connected via bus **380**. The database storage system **140** can store one or more databases the at least one memory **340**, which can include a medical scan database **342** that includes a plurality of medical scan entries **352**, a user database **344** that includes a plurality of user profile entries **354**, a medical scan analysis function database **346** that includes a plurality of medical scan analysis function entries **356**, an interface feature database **348** can include a plurality of interface feature entries **358**, and/or other databases that store data generated and/or utilized by the subsystems **101**. Some or all of the databases **342, 344, 346** and/or **348** can consist of multiple databases, can be stored relationally or non-relationally, and can include different types of entries and different mappings than those described herein. A database entry can include an entry in a relational table or entry in a non-relational structure. Some or all of the data attributes of an entry **352, 354, 356,** and/or **358** can refer to data included in the entry itself or that is otherwise mapped to an identifier included in the entry and can be retrieved from, added to, modified, or deleted from the database storage

system **140** based on a given identifier of the entry. Some or all of the databases **342, 344, 346,** and/or **348** can instead be stored locally by a corresponding subsystem, for example, if they are utilized by only one subsystem.

The processing device **330** can facilitate read/write requests received from subsystems and/or client devices via the network **150** based on read/write permissions for each database stored in the at least one memory device **340**. Different subsystems can be assigned different read/write permissions for each database based on the functions of the subsystem, and different client devices **120** can be assigned different read/write permissions for each database. One or more client devices **120** can correspond to one or more administrators of one or more of the databases stored by the database storage system, and database administrator devices can manage one or more assigned databases, supervise assess and/or efficiency, edit permissions, or otherwise oversee database processes based on input to the client device via interactive interface **275**.

FIG. **4A** presents an embodiment of a medical scan entry **352**, stored in medical scan database **342**, included in metadata of a medical scan, and/or otherwise associated with a medical scan. A medical scan can include imaging data corresponding to a CT scan, x-ray, MM, PET scan, Ultrasound, EEG, mammogram, or other type of radiological scan or medical scan taken of an anatomical region of a human body, animal, organism, or object and further can include metadata corresponding to the imaging data. Some or all of the medical scan entries can be formatted in accordance with a Digital Imaging and Communications in Medicine (DICOM) format or other standardized image format, and some or more of the fields of the medical scan entry **352** can be included in a DICOM header or other standardized header of the medical scan. Medical scans can be awaiting review or can have already been reviewed by one or more users or automatic processes and can include tentative diagnosis data automatically generated by a subsystem, generated based on user input, and/or generated from another source. Some medical scans can include final, known diagnosis data generated by a subsystem and/or generated based on user input, and/or generated from another source, and can included in training sets used to train processes used by one or more subsystems such as the medical scan image analysis system **112** and/or the medical scan natural language analysis system **114**.

Some medical scans can include one or more abnormalities, which can be identified by a user or can be identified automatically. Abnormalities can include nodules, for example malignant nodules identified in a chest CT scan. Abnormalities can also include and/or be characterized by one or more abnormality pattern categories such as such as cardiomegaly, consolidation, effusion, emphysema, and/or fracture, for example identified in a chest x-ray. Abnormalities can also include any other unknown, malignant or benign feature of a medical scan identified as not normal. Some scans can contain zero abnormalities, and can be identified as normal scans. Some scans identified as normal scans can include identified abnormalities that are classified as benign, and include zero abnormalities classified as either unknown or malignant. Scans identified as normal scans may include abnormalities that were not detected by one or more subsystems and/or by an originating entity. Thus, some scans may be improperly identified as normal. Similarly, scans identified to include at least one abnormality may include at least one abnormality that was improperly detected as an abnormality by one or more subsystems

and/or by an originating entity. Thus, some scans may be improperly identified as containing abnormalities.

Each medical scan entry **352** can be identified by its own medical scan identifier **353**, and can include or otherwise map to medical scan image data **410**, and metadata such as scan classifier data **420**, patient history data **430**, diagnosis data **440**, annotation author data **450**, confidence score data **460**, display parameter data **470**, similar scan data **480**, training set data **490**, and/or other data relating to the medical scan. Some or all of the data included in a medical scan entry **352** can be used to aid a user in generating or editing diagnosis data **440**, for example, in conjunction with the medical scan assisted review system **102**, the medical scan report labeling system **104**, and/or the medical scan annotator system **106**. Some or all of the data included in a medical scan entry **352** can be used to allow one or more subsystems **101**, such as automated portions of the medical scan report labeling system **104** and/or the medical scan diagnosing system **108**, to automatically generate and/or edit diagnosis data **440** or other data the medical scan. Some or all of the data included in a medical scan entry **352** can be used to train some or all medical scan analysis functions of the medical scan analysis function database **346** such as one or more medical scan image analysis functions, one or more medical scan natural language analysis functions, one or more medical scan similarity analysis functions, one or more medical report generator functions, and/or one or more medical report analysis functions, for example, in conjunction with the medical scan image analysis system **112**, the medical scan natural language analysis system **114**, and/or the medical scan comparison system **116**.

The medical scan entries **352** and the associated data as described herein can also refer to data associated with a medical scan that is not stored by the medical scan database, for example, that is uploaded by a client device for direct transmission to a subsystem, data generated by a subsystem and used as input to another subsystem or transmitted directly to a client device, data stored by a Picture Archive and Communication System (PACS) communicating with the medical scan processing system **100**, or other data associated with a medical scan that is received and or generated without being stored in the medical scan database **342**. For example, some or all of the structure and data attributes described with respect to a medical scan entry **352** can also correspond to structure and/or data attribute of data objects or other data generated by and/or transmitted between subsystems and/or client devices that correspond to a medical scan. Herein, any of the data attributes described with respect to a medical scan entry **352** can also correspond to data extracted from a data object generated by a subsystem or client device or data otherwise received from a subsystem, client device, or other source via network **150** that corresponds to a medical scan.

The medical scan image data **410** can include one or more images corresponding to a medical scan. The medical scan image data **410** can include one or more image slices **412**, for example, corresponding to a single x-ray image, a plurality of cross-sectional, tomographic images of a scan such as a CT scan, or any plurality of images taken from the same or different point at the same or different angles. The medical scan image data **410** can also indicate an ordering of the one or more image slices **412**. Herein, a “medical scan” can refer to a full scan of any type represented by medical scan image data **410**. Herein, an “image slice” can refer to one of a plurality of cross-sectional images of the medical scan image data **410**, one of a plurality of images taken from different angles of the medical scan image data

410, and/or the single image of the medical scan image data **410** that includes only one image. Furthermore “plurality of image slices” can refer to all of the images of the associated medical scan, and refers to only a single image if the medical scan image data **410** includes only one image. Each image slice **412** can include a plurality of pixel values **414** mapped to each pixel of the image slice. Each pixel value can correspond to a density value, such as a Hounsfield value or other measure of density. Pixel values can also correspond to a grayscale value, a RGB (Red-Green-Blue) or other color value, or other data stored by each pixel of an image slice **412**.

Scan classifier data **420** can indicate classifying data of the medical scan. Scan classifier data can include scan type data **421**, for example, indicating the modality of the scan. The scan classifier data can indicate that the scan is a CT scan, x-ray, MM, PET scan, Ultrasound, EEG, mammogram, or other type of scan. Scan classifier data **420** can also include anatomical region data **422**, indicating for example, the scan is a scan of the chest, head, right knee, or other anatomical region. Scan classifier data can also include originating entity data **423**, indicating the hospital where the scan was taken and/or a user that uploaded the scan to the system. If the originating entity data corresponds to a user of one or more subsystems **101**, the originating entity data can include a corresponding user profile identifier and/or include other data from the user profile entry **354** of the user. Scan classifier data **420** can include geographic region data **424**, indicating a city, state, and/or country from which the scan originated, for example, based on the user data retrieved from the user database **344** based on the originating entity. Scan classifier data can also include machine data **425**, which can include machine identifier data, machine model data, machine calibration data, and/or contrast agent data, for example based on imaging machine data retrieved from the user database **344** based on the originating entity data **423**. The scan classifier data **420** can include scan date data **426** indicating when the scan was taken. The scan classifier data **420** can include scan priority data **427**, which can indicate a priority score, ranking, number in a queue, or other priority data with regard to triaging and/or review. A priority score, ranking, or queue number of the scan priority data **427** can be generated by automatically by a subsystem based on the scan priority data **427**, based on a severity of patient symptoms or other indicators in the risk factor data **432**, based on a priority corresponding to the originating entity, based on previously generated diagnosis data **440** for the scan, and/or can be assigned by the originating entity and/or a user of the system.

The scan classifier data **420** can include other classifying data not pictured in FIG. 4A. For example, a set of scans can include medical scan image data **410** corresponding to different imaging planes. The scan classifier data can further include imaging plane data indicating one or more imaging planes corresponding to the image data. For example, the imaging plane data can indicate the scan corresponds to the axial plane, sagittal plane, or coronal plane. A single medical scan entry **352** can include medical scan image data **410** corresponding multiple planes, and each of these planes can be tagged appropriately in the image data. In other embodiments, medical scan image data **410** corresponding to each plane can be stored as separate medical scan entries **352**, for example, with a common identifier indicating these entries belong to the same set of scans.

Alternatively or in addition, the scan classifier data **420** can include sequencing data. For example, a set of scans can include medical scan image data **410** corresponding to

13

different sequences. The scan classifier data can further include sequencing data indicating one or more of a plurality of sequences of the image data corresponds to, for example, indicating whether an MRI scan corresponds to a T2 sequence, a T1 sequence, a T1 sequence with contrast, a diffusion sequence, a FLAIR sequence, or other MM sequence. A single medical scan entry **352** can include medical scan image data **410** corresponding to multiple sequences, and each of these sequences can be tagged appropriately in the entry. In other embodiments, medical scan image data **410** corresponding to each sequence can be stored as separate medical scan entries **352**, for example, with a common identifier indicating these entries belong to the same set of scans.

Alternatively or in addition, the scan classifier data **420** can include an image quality score. This score can be determined automatically by one or more subsystems **101**, and/or can be manually assigned the medical scan. The image quality score can be based on a resolution of the image data **410**, where higher resolution image data is assigned a more favorable image quality score than lower resolution image data. The image quality score can be based on whether the image data **410** corresponds to digitized image data received directly from the corresponding imaging machine, or corresponds to a hard copy of the image data that was later scanned in. In some embodiments, the image quality score can be based on a detected corruption, and/or detected external factor that determined to negatively affect the quality of the image data during the capturing of the medical scan and/or subsequent to the capturing of the medical scan. In some embodiments, the image quality score can be based on detected noise in the image data, where a medical scan with a higher level of detected noise can receive a less favorable image quality score than a medical scan with a lower level of detected noise. Medical scans with this determined corruption or external factor can receive a less favorable image quality score than medical scans with no detected corruption or external factor.

In some embodiments, the image quality score can be based on include machine data **425**. In some embodiments, one or more subsystems can utilize the image quality score to flag medical scans with image quality scores that fall below an image quality threshold. The image quality threshold can be the same or different for different subsystems, medical scan modalities, and/or anatomical regions. For example, the medical scan image analysis system can automatically filter training sets based on selecting only medical scans with image quality scores that compare favorably to the image quality threshold. As another example, one or more subsystems can flag a particular imaging machine and/or hospital or other medical entity that have produced at least a threshold number and/or percentage of medical scan with image quality scores that compare unfavorably to the image quality threshold. As another example, a de-noising algorithm can be automatically utilized to clean the image data when the image quality score compares unfavorably to the image quality threshold. As another example, the medical scan image analysis system can select a particular medical image analysis function from a set of medical image analysis functions to utilize on a medical scan to generate inference data for the medical scan. Each of this set of medical image analysis function can be trained on different levels of image quality, and the selected image analysis function can be selected based on the determined image quality score falling within a range of image quality scores the image analysis function was trained on and/or is otherwise suitable for.

14

The patient history data **430** can include patient identifier data **431** which can include basic patient information such as name or an identifier that may be anonymized to protect the confidentiality of the patient, age, and/or gender. The patient identifier data **431** can also map to a patient entry in a separate patient database stored by the database storage system, or stored elsewhere. The patient history data can include patient risk factor data **432** which can include previous medical history, family medical history, smoking and/or drug habits, pack years corresponding to tobacco use, environmental exposures, patient symptoms, etc. The patient history data **430** can also include longitudinal data **433**, which can identify one or more additional medical scans corresponding to the patient, for example, retrieved based on patient identifier data **431** or otherwise mapped to the patient identifier data **431**. Some or all additional medical scans can be included in the medical scan database, and can be identified based on their corresponding identifiers medical scan identifiers **353**. Some or all additional medical scans can be received from a different source and can otherwise be identified. Alternatively or in addition, the longitudinal data can simply include some or all relevant scan entry data of a medical scan entry **352** corresponding to the one or more additional medical scans. The additional medical scans can be the same type of scan or different types of scans. Some or all of the additional scans may correspond to past medical scans, and/or some or all of the additional scans may correspond to future medical scans. The longitudinal data **433** can also include data received and/or determined at a date after the scan such as final biopsy data, or some or all of the diagnosis data **440**. The patient history data can also include a longitudinal quality score **434**, which can be calculated automatically by a subsystem, for example, based on the number of additional medical scans, based on how many of the additional scans in the file were taken before and/or after the scan based on the scan date data **426** of the medical scan and the additional medical scans, based on a date range corresponding to the earliest scan and corresponding to the latest scan, based on the scan types data **421** these scans, and/or based on whether or not a biopsy or other final data is included. As used herein, a “high” longitudinal quality score refers to a scan having more favorable longitudinal data than that with a “low” longitudinal quality score.

Diagnosis data **440** can include data that indicates an automated diagnosis, a tentative diagnosis, and/or data that can otherwise be used to support medical diagnosis, triage, medical evaluation and/or other review by a medical professional or other user. The diagnosis data **440** of a medical scan can include a binary abnormality identifier **441** indicating whether the scan is normal or includes at least one abnormality. In some embodiments, the binary abnormality identifier **441** can be determined by comparing some or all of confidence score data **460** to a threshold, can be determined by comparing a probability value to a threshold, and/or can be determined by comparing another continuous or discrete value indicating a calculated likelihood that the scan contains one or more abnormalities to a threshold. In some embodiments, non-binary values, such as one or more continuous or discrete values indicating a likelihood that the scan contains one or more abnormalities, can be included in diagnosis data **440** in addition to, or instead of, binary abnormality identifier **441**. One or more abnormalities can be identified by the diagnosis data **440**, and each identified abnormality can include its own set of abnormality annotation data **442**. Alternatively, some or all of the diagnosis data **440** can indicate and/or describe multiple abnormalities, and

thus will not be presented for each abnormality in the abnormality annotation data 442. For example, the report data 449 of the diagnosis data 440 can describe all identified abnormalities, and thus a single report can be included in the diagnosis.

FIG. 4B presents an embodiment of the abnormality annotation data 442. The abnormality annotation data 442 for each abnormality can include abnormality location data 443, which can include an anatomical location and/or a location specific to pixels, image slices, coordinates or other location information identifying regions of the medical scan itself. The abnormality annotation data 442 can include abnormality classification data 445 which can include binary, quantitative, and/or descriptive data of the abnormality as a whole, or can correspond to one or more abnormality classifier categories 444, which can include size, volume, pre-post contrast, doubling time, calcification, components, smoothness, spiculation, lobulation, sphericity, internal structure, texture, or other categories that can classify and/or otherwise characterize an abnormality. Abnormality classifier categories 444 can be assigned a binary value, indicating whether or not such a category is present. For example, this binary value can be determined by comparing some or all of confidence score data 460 to a threshold, can be determined by comparing a probability value to a threshold, and/or can be determined by comparing another continuous or discrete value indicating a calculated likelihood that a corresponding abnormality classifier category 444 is present to a threshold, which can be the same or different threshold for each abnormality classifier category 444. In some embodiments, abnormality classifier categories 444 can be assigned one or more non-binary values, such as one or more continuous or discrete values indicating a likelihood that the corresponding classifier category 444 is present.

The abnormality classifier categories 444 can also include a malignancy category, and the abnormality classification data 445 can include a malignancy rating such as a Lung-RADS score, a Fleischner score, and/or one or more calculated values that indicate malignancy level, malignancy severity, and/or probability of malignancy. Alternatively or in addition, the malignancy category can be assigned a value of “yes”, “no”, or “maybe”. The abnormality classifier categories 444 can also include abnormality pattern categories 446 such as cardiomegaly, consolidation, effusion, emphysema, and/or fracture, and the abnormality classification data 445 for each abnormality pattern category 446 can indicate whether or not each of the abnormality patterns is present.

The abnormality classifier categories can correspond to Response Evaluation Criteria in Solid Tumors (RECIST) eligibility and/or RECIST evaluation categories. For example, an abnormality classifier category 444 corresponding to RECIST eligibility can have corresponding abnormality classification data 445 indicating a binary value “yes” or “no”, and/or can indicate if the abnormality is a “target lesion” and/or a “non-target lesion.” As another example, an abnormality classifier category 444 corresponding to a RECIST evaluation category can be determined based on longitudinal data 433 and can have corresponding abnormality classification data 445 that includes one of the set of possible values “Complete Response”, “Partial Response”, “Stable Disease”, or “Progressive Disease.”

The diagnosis data 440 as a whole, and/or the abnormality annotation data 442 for each abnormality, can include custom codes or datatypes identifying the binary abnormality identifier 441, abnormality location data 443 and/or some or

all of the abnormality classification data 445 of one or more abnormality classifier categories 444. Alternatively or in addition, some or all of the abnormality annotation data 442 for each abnormality and/or other diagnosis data 440 can be presented in a DICOM format or other standardized image annotation format, and/or can be extracted into custom datatypes based on abnormality annotation data originally presented in DICOM format. Alternatively or in addition, the diagnosis data 440 and/or the abnormality annotation data 442 for each abnormality can be presented as one or more medical codes 447 such as SNOMED codes, Current Procedure Technology (CPT) codes, ICD-9 codes, ICD-10 codes, or other standardized medical codes used to label or otherwise describe medical scans.

Alternatively or in addition, the diagnosis data 440 can include natural language text data 448 annotating or otherwise describing the medical scan as a whole, and/or the abnormality annotation data 442 can include natural language text data 448 annotating or otherwise describing each corresponding abnormality. In some embodiments, some or all of the diagnosis data 440 is presented only as natural language text data 448. In some embodiments, some or all of the diagnosis data 440 is automatically generated by one or more subsystems based on the natural language text data 448, for example, without utilizing the medical scan image data 410, for example, by utilizing one or more medical scan natural language analysis functions trained by the medical scan natural language analysis system 114. Alternatively or in addition, some embodiments, some or all of the natural language text data 448 is generated automatically based on other diagnosis data 440 such as abnormality annotation data 442, for example, by utilizing a medical scan natural language generating function trained by the medical scan natural language analysis system 114.

The diagnosis data can include report data 449 that includes at least one medical report, which can be formatted to include some or all of the medical codes 447, some or all of the natural language text data 448, other diagnosis data 440, full or cropped images slices formatted based on the display parameter data 470 and/or links thereto, full or cropped images slices or other data based on similar scans of the similar scan data 480 and/or links thereto, full or cropped images or other data based on patient history data 430 such as longitudinal data 433 and/or links thereto, and/or other data or links to data describing the medical scan and associated abnormalities. The diagnosis data 440 can also include finalized diagnosis data corresponding to future scans and/or future diagnosis for the patient, for example, biopsy data or other longitudinal data 433 determined subsequently after the scan. The medical report of report data 449 can be formatted based on specified formatting parameters such as font, text size, header data, bulleting or numbering type, margins, file type, preferences for including one or more full or cropped image slices 412, preferences for including similar medical scans, preferences for including additional medical scans, or other formatting to list natural language text data and/or image data, for example, based on preferences of a user indicated in the originating entity data 423 or other responsible user in the corresponding report formatting data.

Annotation author data 450 can be mapped to the diagnosis data for each abnormality, and/or mapped to the scan as a whole. This can include one or more annotation author identifiers 451, which can include one or more user profile identifiers of a user of the system, such as an individual medical professional, medical facility and/or medical entity that uses the system. Annotation author data 450 can be used

to determine the usage data of a user profile entry **354**. Annotation author data **450** can also include one or more medical scan analysis function identifiers **357** or other function identifier indicating one or more functions or other processes of a subsystem responsible for automatically generating and/or assisting a user in generating some or all of the diagnosis data, for example an identifier of a particular type and/or version of a medical scan image analysis functions that was used by the medical scan diagnosing system **108** used to generate part or all of the diagnosis data **440** and/or an interface feature identifier, indicating an one or more interface features presented to a user to facilitate entry of and/or reviewing of the diagnosis data **440**. The annotation author data can also simply indicate, for one or more portions of the diagnosis data **440**, if this portion was generated by a human or automatically generated by a subsystem of the medical scan processing system.

In some embodiments, if a medical scan was reviewed by multiple entities, multiple, separate diagnosis data entries **440** can be included in the medical scan entry **352**, mapped to each diagnosis author in the annotation author data **450**. This allows different versions of diagnosis data **440** received from multiple entities. For example, annotation author data of a particular medical scan could indicate that the annotation data was written by a doctor at medical entity A, and the medical code data was generated by user Y by utilizing the medical scan report labeling system **104**, which was confirmed by expert user X. The annotation author data of another medical scan could indicate that the medical code was generated automatically by utilizing version 7 of the medical scan image analysis function relating to chest x-rays, and confirmed by expert user X. The annotation author data of another medical scan could indicate that the location and a first malignancy rating were generated automatically by utilizing version 7 of the medical scan image analysis function relating to chest x-rays, and that a second malignancy rating was entered by user Z. In some embodiments, one of the multiple diagnosis entries can include consensus annotation data, for example, generated automatically by a subsystem such as the medical scan annotating system **106** based on the multiple diagnosis data **440**, based on confidence score data **460** of each of the multiple diagnosis data **440**, and/or based on performance score data of a corresponding user, a medical scan analysis function, or an interface feature, identified in the annotation author data for each corresponding one of the multiple diagnosis data **440**.

Confidence score data **460** can be mapped to some or all of the diagnosis data **440** for each abnormality, and/or for the scan as a whole. This can include an overall confidence score for the diagnosis, a confidence score for the binary indicator of whether or not the scan was normal, a confidence score for the location a detected abnormality, and/or confidence scores for some or all of the abnormality classifier data. This may be generated automatically by a subsystem, for example, based on the annotation author data and corresponding performance score of one or more identified users and/or subsystem attributes such as interactive interface types or medical scan image analysis functions indicated by the annotation author data. In the case where multiple diagnosis data entries **440** are included from different sources, confidence score data **460** can be computed for each entry and/or an overall confidence score, for example, corresponding to consensus diagnosis data, can be based on calculated distance or other error and/or discrepancies between the entries, and/or can be weighted on the confidence score data **460** of each entry. In various embodiments,

the confidence score data **460** can include a truth flag **461** indicating the diagnosis data is considered as “known” or “truth”, for example, flagged based on user input, flagged automatically based on the author data, and/or flagged automatically based on the calculated confidence score of the confidence score data exceeding a truth threshold. As used herein, a “high” confidence score refers to a greater degree or more favorable level of confidence than a “low” confidence score.

Display parameter data **470** can indicate parameters indicating an optimal or preferred display of the medical scan by an interactive interface **275** and/or formatted report for each abnormality and/or for the scan as a whole. Some or all of the display parameter data can have separate entries for each abnormality, for example, generated automatically by a subsystem **101** based on the abnormality annotation data **442**. Display parameter data **470** can include interactive interface feature data **471**, which can indicate one or more selected interface features associated with the display of abnormalities and/or display of the medical scan as a whole, and/or selected interface features associated with user interaction with a medical scan, for example, based on categorized interface feature performance score data and a category associated with the abnormality and/or with the medical scan itself. The display parameter data can include a slice subset **472**, which can indicate a selected subset of the plurality of image slices that includes a single image slice **412** or multiple image slices **412** of the medical scan image data **410** for display by a user interface. The display parameter data **470** can include slice order data **473** that indicates a selected custom ordering and/or ranking for the slice subset **472**, or for all of the slices **412** of the medical scan. The display parameter data **470** can include slice cropping data **474** corresponding to some or all of the slice subset **472**, or all of the image slices **412** of the medical scan, and can indicate a selected custom cropped region of each image slice **412** for display, or the same selected custom cropped region for the slice subset **472** or for all slices **412**. The display parameter data can include density window data **475**, which can indicate a selected custom density window for display of the medical scan as a whole, a selected custom density window for the slice subset **472**, and/or selected custom density windows for each of the image slices **412** of the slice subset **472**, and/or for each image slice **412** of the medical scan. The density window data **475** can indicate a selected upper density value cut off and a selected lower density value cut off, and/or can include a selected deterministic function to map each density value of a pixel to a grayscale value based on the preferred density window. The interactive interface feature data **471**, slice subset **472**, slice order data **473**, slice cropping data **474**, and/or the density window data **475** can be selected via user input and/or generated automatically by one or more subsystems **101**, for example, based on the abnormality annotation data **442** and/or based on performance score data of different interactive interface versions.

Similar scan data **480** can be mapped to each abnormality, or the scan as a whole, and can include similar scan identifier data **481** corresponding to one or more identified similar medical scans, for example, automatically identified by a subsystem **101**, for example, by applying a similar scan identification step of the medical scan image analysis system **112** and/or applying medical scan similarity analysis function to some or all of the data stored in the medical scan entry of the medical scan, and/or to some or all corresponding data of other medical scans in the medical scan database. The similar scan data **480** can also correspond to medical

scans received from another source. The stored similarity data can be used to present similar cases to users of the system and/or can be used to train medical scan image analysis functions or medical scan similarity analysis functions.

Each identified similar medical scan can have its own medical scan entry **352** in the medical scan database **342** with its own data, and the similar scan identifier data **481** can include the medical scan identifier **353** each similar medical scan. Each identified similar medical scan can be a scan of the same scan type or different scan type than medical scan.

The similar scan data **480** can include a similarity score **482** for each identified similar scan, for example, generated based on some or all of the data of the medical scan entry **352** for medical scan and based on some or all of the corresponding data of the medical scan entry **352** for the identified similar medical scan. For example, the similarity score **482** can be generated based on applying a medical scan similarity analysis function to the medical image scan data of medical scans and **402**, to some or all of the abnormality annotation data of medical scans and **402**, and/or to some or all of the patient history data **430** of medical scans and **402** such as risk factor data **432**. As used herein, a “high” similarity score refers a higher level of similarity that a “low” similarity score.

The similar scan data **480** can include its own similar scan display parameter data **483**, which can be determined based on some or all of the display parameter data **470** of the identified similar medical scan. Some or all of the similar scan display parameter data **483** can be generated automatically by a subsystem, for example, based on the display parameter data **470** of the identified similar medical scan, based on the abnormality annotation data **442** of the medical scan itself and/or based on display parameter data **470** of the medical scan itself. Thus, the similar scan display parameter data **483** can be the same or different than the display parameter data **470** mapped to the identified similar medical scan and/or can be the same or different than the display parameter data **470** of the medical scan itself. This can be utilized when displaying similar scans to a user via interactive interface **275** and/or can be utilized when generating report data **449** that includes similar scans, for example, in conjunction with the medical scan assisted review system **102**.

The similar scan data **480** can include similar scan abnormality data **484**, which can indicate one of a plurality of abnormalities of the identified similar medical scan and its corresponding abnormality annotation data **442**. For example, the similarity scan abnormality data **484** can include an abnormality pair that indicates one of a plurality of abnormalities of the medical scan, and indicates one of a plurality of abnormalities of the identified similar medical scan, for example, that was identified as the similar abnormality.

The similar scan data **480** can include similar scan filter data **485**. The similar scan filter data can be generated automatically by a subsystem, and can include a selected ordered or un-ordered subset of all identified similar scans of the similar scan data **480**, and/or a ranking of all identified similar scans. For example, the subset can be selected and/or some or all identified similar scans can be ranked based on each similarity score **482**, and/or based on other factors such as based on a longitudinal quality score **434** of each identified similar medical scan.

The training set data **490** can indicate one or more training sets that the medical scan belongs to. For example, the training set data can indicate one or more training set

identifiers **491** indicating one or more medical scan analysis functions that utilized the medical scan in their training set, and/or indicating a particular version identifier **641** of the one or more medical scan analysis functions that utilized the medical scan in their training set. The training set data **490** can also indicate which portions of the medical scan entry were utilized by the training set, for example, based on model parameter data **623** of the corresponding medical scan analysis functions. For example, the training set data **490** can indicate that the medical scan image data **410** was included in the training set utilized to train version X of the chest x-ray medical scan image analysis function, or that the natural language text data **448** of this medical scan was used to train version Y of the natural language analysis function.

FIG. 5A presents an embodiment of a user profile entry **354**, stored in user database **344** or otherwise associated with a user. A user can correspond to a user of one or more of the subsystems such as a radiologist, doctor, medical professional, medical report labeler, administrator of one or more subsystems or databases, or other user that uses one or more subsystems **101**. A user can also correspond to a medical entity such as a hospital, medical clinic, establishment that utilizes medical scans, establishment that employs one or more of the medical professionals described, an establishment associated with administering one or more subsystems, or other entity. A user can also correspond to a particular client device **120** or account that can be accessed one or more medical professionals or other employees at the same or different medical entities. Each user profile entry can have a corresponding user profile identifier **355**.

A user profile entry **354** can include basic user data **510**, which can include identifying information **511** corresponding to the user such as a name, contact information, account/login/password information, geographic location information such as geographic region data **424**, and/or other basic information. Basic user data **510** can include affiliation data **512**, which can list one or more medical entities or other establishments the user is affiliated with, for example, if the user corresponds to a single person such as a medical professional, or if the user corresponds to a hospital in a network of hospitals. The affiliation data **512** can include one or more corresponding user profile identifiers **355** and/or basic user data **510** if the corresponding affiliated medical entity or other establishment has its own entry in the user database. The user identifier data can include employee data **513** listing one or more employees, such as medical professionals with their own user profile entries **354**, for example, if the user corresponds to a medical entity or supervising medical professional of other medical professional employees, and can list a user profile identifier **355** and/or basic user data **510** for each employee. The basic user data **510** can also include imaging machine data **514**, which can include a list of machines affiliated with the user which can include machine identifiers, model information, calibration information, scan type information, or other data corresponding to each machine, for example, corresponding to the machine data **425**. The user profile entry can include client device data **515**, which can include identifiers for one or more client devices associated with the user, for example, allowing subsystems **101** to send data to a client device **120** corresponding to a selected user based on the client device data and/or to determine a user that data was received by determining the client device from which the data was received.

The user profile entry can include usage data **520** which can include identifying information for a plurality of usages by the user in conjunction with using one or more subsystems **101**. This can include consumption usage data **521**,

21

which can include a listing of, or aggregate data associated with, usages of one or more subsystems by the user, for example, where the user is utilizing the subsystem as a service. For example, the consumption usage data **521** can correspond to each instance where diagnosis data was sent to the user for medical scans provided to the user in conjunction with the medical scan diagnosing system **108** and/or the medical scan assisted review system **102**. Some or all of consumption usage data **521** can include training usage data **522**, corresponding to usage in conjunction with a certification program or other user training provided by one or more subsystems. The training usage data **522** can correspond to each instance where diagnosis feedback data was provided by user for a medical scan with known diagnosis data, but diagnosis feedback data is not utilized by a subsystem to generate, edit, and/or confirm diagnosis data **440** of the medical scan, as it is instead utilized to train a user and/or determine performance data for a user.

Usage data **520** can include contribution usage data **523**, which can include a listing of, or aggregate data associated with, usages of one or more subsystems **101** by the user, for example, where the user is generating and/or otherwise providing data and/or feedback that can be utilized by the subsystems, for example, to generate, edit, and/or confirm diagnosis data **440** and/or to otherwise populate, modify, or confirm portions of the medical scan database or other subsystem data. For example, the contribution usage data **523** can correspond to diagnosis feedback data received from user, used to generate, edit, and/or confirm diagnosis data. The contribution usage data **523** can include interactive interface feature data **524** corresponding to the interactive interface features utilized with respect to the contribution.

The consumption usage data **521** and/or the contribution usage data **523** can include medical scan entry **352** whose entries the user utilized and/or contributed to, can indicate one or more specific attributes of a medical scan entry **352** that a user utilized and/or contributed to, and/or a log of the user input generated by a client device of the user in conjunction with the data usage. The contribution usage data **523** can include the diagnosis data that the user may have generated and/or reviewed, for example, indicated by, mapped to, and/or used to generate the annotation author data **450** of corresponding medical scan entries **352**. Some usages may correspond to both consumption usage of the consumption usage data **521** and contribution usage of the contribution usage data **523**. The usage data **520** can also indicate one or more subsystems **101** that correspond to each consumption and/or contribution.

The user profile entry can include performance score data **530**. This can include one or more performance scores generated based on the contribution usage data **523** and/or training usage data **522**. The performance scores can include separate performance scores generated for every contribution in the contribution usage data **523** and/or training usage data **522** and/or generated for every training consumption usages corresponding to a training program. As used herein, a "high" performance score refers to a more favorable performance or rating than a "low" performance score.

The performance score data can include accuracy score data **531**, which can be generated automatically by a subsystem for each contribution, for example, based on comparing diagnosis data received from a user to data to known truth data such as medical scans with a truth flag **461**, for example, retrieved from the corresponding medical scan entry **352** and/or based on other data corresponding to the medical scan, for example, received from an expert user that later reviewed the contribution usage data of the user and/or

22

generated automatically by a subsystem. The accuracy score data **531** can include an aggregate accuracy score generated automatically by a subsystem, for example, based on the accuracy data of multiple contributions by the user over time.

The performance data can also include efficiency score data **532** generated automatically by a subsystem for each contribution based on an amount of time taken to complete a contribution, for example, from a time the request for a contribution was sent to the client device to a time that the contribution was received from the client device, based on timing data received from the client device itself, and/or based on other factors. The efficiency score can include an aggregate efficiency score, which can be generated automatically by a subsystem based on the individual efficiency scores over time and/or based on determining a contribution completion rate, for example based on determining how many contributions were completed in a fixed time window.

Aggregate performance score data **533** can be generated automatically by a subsystem based on the aggregate efficiency and/or accuracy data. The aggregate performance data can include categorized performance data **534**, for example, corresponding to different scan types, different anatomical regions, different subsystems, different interactive interface features and/or display parameters. The categorized performance data **534** can be determined automatically by a subsystem based on the scan type data **421** and/or anatomical region data **422** of the medical scan associated with each contribution, one or more subsystems **101** associated with each contribution, and/or interactive interface feature data **524** associated with each contribution. The aggregate performance data can also be based on performance score data **530** of individual employees if the user corresponds to a medical entity, for example, retrieved based on user profile identifiers **355** included in the employee data **513**. The performance score data can also include ranking data **535**, which can include an overall ranking or categorized rankings, for example, generated automatically by a subsystem or the database itself based on the aggregate performance data.

In some embodiments, aggregate data for each user can be further broken down based on scores for distinct scan categories, for example, based on the scan classifier data **420**, for example, where a first aggregate data score is generated for a user "A" based on scores from all knee x-rays, and a second aggregate data score is generated for user A based on scores from all chest CT scans. Aggregate data for each user can be further based on scores for distinct diagnosis categories, where a first aggregate data score is generated for user A based on scores from all normal scans, and a second aggregate data score is generated for user A based on scores from all scans that contain an abnormality. This can be further broken down, where a first aggregate score is generated for user A based on all scores from scans that contain an abnormality of a first type and/or in a first anatomical location, and a second aggregate score is generated for A based on all scores from scans that contain an abnormality of a second type and/or in a second location. Aggregate data for each user can be further based on affiliation data, where a ranking is generated for a medical professional "B" based on scores from all medical professionals with the same affiliation data, and/or where a ranking is generated for a hospital "C" based on scores for all hospitals, all hospitals in the same geographical region, etc. Aggregate data for each user can be further based on scores for interface features, where a first aggregate data score is generated for user A based on scores using a first interface

feature, and a second aggregate data score is generated for user A based on scores using a first interface feature.

The user profile entry can include qualification data **540**. The qualification data can include experience data **541** such as education data, professional practice data, number of years practicing, awards received, etc. The qualification data **540** can also include certification data **542** corresponding to certifications earned based on contributions to one or more subsystems, for example, assigned to users automatically by a subsystem based on the performance score data **530** and/or based on a number of contributions in the contribution usage data **523** and/or training usage data **522**. For example, the certifications can correspond to standard and/or recognized certifications to train medical professionals and/or incentivize medical professionals to use the system. The qualification data **540** can include expert data **543**. The expert data **543** can include a binary expert identifier, which can be generated automatically by a subsystem based on experience data **541**, certification data **542**, and/or the performance score data **530**, and can indicate whether the user is an expert user. The expert data **543** can include a plurality of categorized binary expert identifiers corresponding to a plurality of qualification categories corresponding to corresponding to scan types, anatomical regions, and/or the particular subsystems. The categorized binary expert identifiers can be generated automatically by a subsystem based on the categorized performance data **534** and/or the experience data **541**. The categories be ranked by performance score in each category to indicate particular specialties. The expert data **543** can also include an expert ranking or categorized expert ranking with respect to all experts in the system.

The user profile entry can include subscription data **550**, which can include a selected one of a plurality of subscription options that the user has subscribed to. For example, the subscription options can correspond to allowed usage of one or more subsystems, such as a number of times a user can utilize a subsystem in a month, and/or to a certification program, for example paid for by a user to receive training to earn a subsystem certification of certification data **542**. The subscription data can include subscription expiration information, and/or billing information. The subscription data can also include subscription status data **551**, which can for example indicate a number of remaining usages of a system and/or available credit information. For example, the remaining number of usages can decrease and/or available credit can decrease in response to usages that utilize one or more subsystems as a service, for example, indicated in the consumption usage data **521** and/or training usage data **522**. In some embodiments, the remaining number of usages can increase and/or available credit can increase in response to usages that correspond to contributions, for example, based on the contribution usage data **523**. An increase in credit can be variable, and can be based on a determined quality of each contribution, for example, based on the performance score data **530** corresponding to the contribution where a higher performance score corresponds to a higher increase in credit, based on scan priority data **427** of the medical scan where contributing to higher priority scans corresponds to a higher increase in credit, or based on other factors.

The user profile entry **354** can include interface preference data **560**. The interface preference data can include a preferred interactive interface feature set **561**, which can include one or more interactive interface feature identifiers and/or one or more interactive interface version identifiers of interface feature entries **358** and/or version identifiers of the interface features. Some or all of the interface features of the preferred interactive interface feature set **561** can correspond

to display parameter data **470** of medical scans. The preferred interactive interface feature set **561** can include a single interactive feature identifier for one or more feature types and/or interface types, and/or can include a single interactive interface version identifier for one or more interface categories. The preferred interactive interface feature set **561** can include a ranking of multiple features for the same feature type and/or interface type. The ranked and/or unranked preferred interactive interface feature set **561** can be generated based on user input to an interactive interface of the client device to select and/or rank some or all of the interface features and/or versions. Some or all of the features and/or versions of the preferred interactive feature set can be selected and/or ranked automatically by a subsystem such as the medical scan interface evaluator system, for example based on interface feature performance score data and/or feature popularity data. Alternatively or in addition, the performance score data **530** can be utilized by a subsystem to automatically determine the preferred interactive feature set, for example, based on the scores in different feature-based categories of the categorized performance data **534**.

The user profile entry **354** can include report formatting data **570**, which can indicate report formatting preferences indicated by the user. This can include font, text size, header data, bulleting or numbering type, margins, file type, preferences for including one or more full or cropped image slices **412**, preferences for including similar medical scans, preferences for including additional medical scans in reports, or other formatting preference to list natural language text data and/or image data corresponding to each abnormality. Some or all of the report formatting data **570** can be based on interface preference data **560**. The report formatting data **570** can be used by one or more subsystems to automatically generate report data **449** of medical scans based on the preferences of the requesting user.

FIG. **5B** presents an embodiment of a medical scan analysis function entry **356**, stored in medical scan analysis function database **346** or otherwise associated with one of a plurality of medical scan analysis functions trained by and/or utilized by one or more subsystems **101**. For example, a medical scan analysis function can include one or more medical scan image analysis functions trained by the medical scan image analysis system **112**; one or more medical scan natural language analysis functions trained by the medical scan natural language analysis system **114**; one or more medical scan similarity analysis function trained by the medical scan image analysis system **112**, the medical scan natural language analysis system **114**, and/or the medical scan comparison system **116**; one or more medical report generator functions trained by the medical scan natural language analysis system **114** and/or the medical scan image analysis system **112**, and/or the medical report analysis function trained by the medical scan natural language analysis system **114**. Some or all of the medical scan analysis functions can correspond to medical scan inference functions of the medical scan diagnosing system **108**, the de-identification function and/or the inference functions utilized by a medical picture archive integration system as discussed in conjunction with FIGS. **8A-8F**, or other functions and/or processes described herein in conjunction with one or more subsystems **101**. Each medical scan analysis function entry **356** can include a medical scan analysis function identifier **357**.

A medical scan analysis function entry **356** can include function classifier data **610**. Function classifier data **610** can include input and output types corresponding to the function. For example the function classifier data can include

input scan category **611** that indicates which types of scans can be used as input to the medical scan analysis function. For example, input scan category **611** can indicate that a medical scan analysis function is for chest CT scans from a particular hospital or other medical entity. The input scan category **611** can include one or more categories included in scan classifier data **420**. In various embodiments, the input scan category **611** corresponds to the types of medical scans that were used to train the medical scan analysis function. Function classifier data **610** can also include output type data **612** that characterizes the type of output that will be produced by the function, for example, indicating that a medical scan analysis function is used to generate medical codes **447**. The input scan category **611** can also include information identifying which subsystems **101** are responsible for running the medical scan analysis function.

A medical scan analysis function entry **356** can include training parameters **620**. This can include training set data **621**, which can include identifiers for the data used to train the medical scan analysis function, such as a set of medical scan identifiers **353** corresponding to the medical scans used to train the medical scan analysis function, a list of medical scan reports and corresponding medical codes used to train the medical scan analysis function, etc. Alternatively or in addition to identifying particular scans of the training set, the training set data **621** can identify training set criteria, such as necessary scan classifier data **420**, necessary abnormality locations, classifiers, or other criteria corresponding to abnormality annotation data **442**, necessary confidence score data **460**, for example, indicating that only medical scans with diagnosis data **440** assigned a truth flag **461** or with confidence score data **460** otherwise comparing favorably to a training set confidence score threshold are included, a number of medical scans to be included and proportion data corresponding to different criteria, or other criteria used to populate a training set with data of medical scans. Training parameters **620** can include model type data **622** indicating one or more types of model, methods, and/or training functions used to determine the medical scan analysis function by utilizing the training set **621**. Training parameters **620** can include model parameter data **623** that can include a set of features of the training data selected to train the medical scan analysis function, determined values for weights corresponding to selected input and output features, determined values for model parameters corresponding to the model itself, etc. The training parameter data can also include testing data **624**, which can identify a test set of medical scans or other data used to test the medical scan analysis function. The test set can be a subset of training set **621**, include completely separate data than training set **621**, and/or overlap with training set **621**. Alternatively or in addition, testing data **624** can include validation parameters such as a percentage of data that will be randomly or pseudo-randomly selected from the training set for testing, parameters characterizing a cross validation process, or other information regarding testing. Training parameters **620** can also include training error data **625** that indicates a training error associated with the medical scan analysis function, for example, based on applying cross validation indicated in testing data **624**.

A medical scan analysis function entry **356** can include performance score data **630**. Performance data can include model accuracy data **631**, for example, generated and/or updated based on the accuracy of the function when performed on new data. For example, the model accuracy data **631** can include or be calculated based on the model error for determined for individual uses, for example, generated by

comparing the output of the medical scan analysis function to corresponding data generated by user input to interactive interface **275** in conjunction with a subsystem **101** and/or generated by comparing the output of the medical scan analysis function to medical scans with a truth flag **461**. The model accuracy data **631** can include aggregate model accuracy data computed based on model error of individual uses of the function over time. The performance score data **630** can also include model efficiency data **632**, which can be generated based on how quickly the medical scan analysis function performs, how much memory is utilized by medical scan analysis function, or other efficiency data relating to the medical scan analysis function. Some or all of the performance score data **630** can be based on training error data **625** or other accuracy and/or efficiency data determined during training and/or validation. As used herein, a "high" performance score refers to a more favorable performance or rating than a "low" performance score.

A medical scan analysis function entry **356** can include version data **640**. The version data can include a version identifier **641**. The version data can indicate one or more previous version identifiers **642**, which can map to version identifiers **641** stored in other medical scan analysis function entry **356** that correspond to previous versions of the function. Alternatively or in addition, the version data can indicate multiple versions of the same type based on function classifier data **610**, can indicate the corresponding order and/or rank of the versions, and/or can indicate training parameters **620** associated with each version.

A medical scan analysis function entry **356** can include remediation data **650**. Remediation data **650** can include remediation instruction data **651** which can indicate the steps in a remediation process indicating how a medical scan analysis function is taken out of commission and/or reverted to a previous version in the case that remediation is necessary. The version data **640** can further include remediation criteria data **652**, which can include threshold data or other criteria used to automatically determine when remediation is necessary. For example, the remediation criteria data **652** can indicate that remediation is necessary at any time where the model accuracy data and/or the model efficiency data compares unfavorably to an indicated model accuracy threshold and/or indicated model efficiency threshold. The remediation data **650** can also include recommissioning instruction data **653**, identifying required criteria for recommissioning a medical scan analysis function and/or updating a medical scan analysis function. The remediation data **650** can also include remediation history, indicating one or more instances that the medical scan analysis function was taken out of commission and/or was recommissioned.

FIGS. **6A** and **6B** present an embodiment of a medical scan diagnosing system **108**.

The medical scan diagnosing system **108** can generate inference data **1110** for medical scans by utilizing a set of medical scan inference functions **1105**, stored and run locally, stored and run by another subsystem **101**, and/or stored in the medical scan analysis function database **346**, where the function and/or parameters of the function can be retrieved from the database by the medical scan diagnosing system. For example, the set of medical scan inference function **1105** can include some or all medical scan analysis functions described herein or other functions that generate inference data **1110** based on some or all data corresponding to a medical scan such as some or all data of a medical scan entry **352**. Each medical scan inference function **1105** in the set can correspond to a scan category **1120**, and can be trained on a set of medical scans that compare favorably to

the scan category **1120**. For example, each inference function can be trained on a set of medical scans of the one or more same scan classifier data **420**, such as the same and/or similar scan types, same and/or similar anatomical regions locations, same and/or similar machine models, same and/or similar machine calibration, same and/or similar contrasting agent used, same and/or similar originating entity, same and/or similar geographical region, and/or other classifiers. Thus, the scan categories **1120** can correspond to one or more of a scan type, scan anatomical region data, hospital or other originating entity data, machine model data, machine calibration data, contrast agent data, geographic region data, and/or other scan classifying data **420**. For example, a first medical scan inference function can be directed to characterizing knee x-rays, and a second medical scan inference function can be directed to chest CT scans. As another example, a first medical scan inference function can be directed to characterizing CT scans from a first hospital, and a second medical scan image analysis function can be directed to characterizing CT scans from a second hospital.

Training on these categorized sets separately can ensure each medical scan inference function **1105** is calibrated according to its scan category **1120**, for example, allowing different inference functions to be calibrated on type specific, anatomical region specific, hospital specific, machine model specific, and/or region-specific tendencies and/or discrepancies. Some or all of the medical scan inference functions **1105** can be trained by the medical scan image analysis system and/or the medical scan natural language processing system, and/or some medical scan inference functions **1105** can utilize both image analysis and natural language analysis techniques to generate inference data **1110**. For example, some or all of the inference functions can utilize image analysis of the medical scan image data **410** and/or natural language data extracted from abnormality annotation data **442** and/or report data **449** as input, and generate diagnosis data **440** such as medical codes **447** as output. Each medical scan inference function can utilize the same or different learning models to train on the same or different features of the medical scan data, with the same or different model parameters, for example indicated in the model type data **622** and model parameter data **623**. Model type and/or parameters can be selected for a particular medical scan inference function based on particular characteristics of the one or more corresponding scan categories **1120**, and some or all of the indicated in the model type data **622** and model parameter data **623** can be selected automatically by a subsystem during the training process based on the particular learned and/or otherwise determined characteristics of the one or more corresponding scan categories **1120**.

As shown in FIG. 6A, the medical scan diagnosing system **108** can automatically select a medical scan for processing in response to receiving it from a medical entity via the network. Alternatively, the medical scan diagnosing system **108** can automatically retrieve a medical scan from the medical scan database that is selected based on a request received from a user for a particular scan and/or based on a queue of scans automatically ordered by the medical scan diagnosing system **108** or another subsystem based on scan priority data **427**.

Once a medical scan to be processed is determined, the medical scan diagnosing system **108** can automatically select an inference function **1105** based on a determined scan category **1120** of the selected medical scan and based on corresponding inference function scan categories. The scan category **1120** of a scan can be determined based one

some or all of the scan classifier data **420** and/or based on other metadata associated with the scan. This can include determining which one of the plurality of medical scan inference functions **1105** matches or otherwise compares favorably to the scan category **1120**, for example, by comparing the scan category **1120** to the input scan category of the function classifier data **610**.

Alternatively or in addition, the medical scan diagnosing system **108** can automatically determine which medical scan inference function **1105** is utilized based on an output preference that corresponding to a desired type of inference data **1110** that is outputted by an inference function **1105**. The output preference designated by a user of the medical scan diagnosing system **108** and/or based on the function of a subsystem **101** utilizing the medical scan diagnosing system **108**. For example, the set of inference functions **1105** can include inference functions that are utilized to indicate whether or not a medical scan is normal, to automatically identify at least one abnormality in the scan, to automatically characterize the at least one abnormality in the scan, to assign one or more medical codes to the scan, to generate natural language text data and/or a formatted report for the scan, and/or to automatically generate other diagnosis data such as some or all of diagnosis data **440** based on the medical scan. Alternatively or in addition, some inference functions can also be utilized to automatically generate confidence score data **460**, display parameter data **470**, and/or similar scan data **480**. The medical scan diagnosing system **108** can compare the output preference to the output type data **612** of the medical scan inference function **1105** to determine the selected inference function **1105**. For example, this can be used to decide between a first medical scan inference function that automatically generates medical codes and a second medical scan inference function that automatically generates natural language text for medical reports based on the desired type of inference data **1110**.

Prior to performing the selected medical scan inference function **1105**, the medical scan diagnosing system **108** can automatically perform an input quality assurance function **1106** to ensure the scan classifier data **420** or other metadata of the medical scan accurately classifies the medical scan such that the appropriate medical scan inference function **1105** of the appropriate scan category **1120** is selected. The input quality assurance function can be trained on, for example, medical scan image data **410** of plurality of previous medical scans with verified scan categories. Thus, the input quality assurance function **1106** can take medical scan image data **410** as input and can generate an inferred scan category as output. The inferred scan category can be compared to the scan category **1120** of the scan, and the input quality assurance function **1106** can determine whether or not the scan category **1120** is appropriate by determining whether the scan category **1120** compares favorably to the automatically generated inferred scan category. The input quality assurance function **1106** can also be utilized to reassign the generated inferred scan category to the scan category **1120** when the scan category **1120** compares favorably to the automatically generated inferred scan category. The input quality assurance function **1106** can also be utilized to assign the generated inferred scan category to the scan category **1120** for incoming medical scans that do not include any classifying data, and/or to add classifiers in scan classifier data **420** to medical scans missing one or more classifiers.

In various embodiments, upon utilizing the input quality assurance function **1106** to determine that the scan category **1120** determined by a scan classifier data **420** or other

metadata is inaccurate, the medical scan diagnosing system **108** can transmit an alert and/or an automatically generated inferred scan category to the medical entity indicating that the scan is incorrectly classified in the scan classifier data **420** or other metadata. In some embodiments, the medical scan diagnosing system **108** can automatically update performance score data corresponding to the originating entity of the scan indicated in originating entity data **423**, or another user or entity responsible for classifying the scan, for example, where a lower performance score is generated in response to determining that the scan was incorrectly classified and/or where a higher performance score is generated in response to determining that the scan was correctly classified.

In some embodiments, the medical scan diagnosing system **108** can transmit the medical scan and/or the automatically generated inferred scan category to a selected user. The user can be presented the medical scan image data **410** and/or other data of the medical scan via the interactive interface **275**, for example, displayed in conjunction with the medical scan assisted review system **102**. The interface can prompt the user to indicate the appropriate scan category **1120** and/or prompt the user to confirm and/or edit the inferred scan category, also presented to the user. For example, scan review data can be automatically generated to reflect the user generated and/or verified scan category **1120**. This user indicated scan category **1120** can be utilized to select to the medical scan inference function **1105** and/or to update the scan classifier data **420** or other metadata accordingly. In some embodiments, for example, where the scan review data indicates that the selected user disagrees with the automatically generated inferred scan category created by the input quality assurance function **1106**, the medical scan diagnosing system **108** can automatically update performance score data **630** of the input quality assurance function **1106** by generating a low performance score and/or determine to enter the remediation step **1140** for the input quality assurance function **1106**.

The medical scan diagnosing system **108** can also automatically perform an output quality assurance step after a medical scan inference function **1105** has been performed on a medical scan to produce the inference data **1110**, as illustrated in the embodiment presented in FIG. **6B**. The output quality assurance step can be utilized to ensure that the selected medical scan inference function **1105** generated appropriate inference data **1110** based on expert feedback. The inference data **1110** generated by performing the selected medical scan inference function **1105** can be sent to a client device **120** of a selected expert user, such as an expert user in the user database selected based on categorized performance data and/or qualification data that corresponds to the scan category **1120** and/or the inference itself, for example, by selecting an expert user best suited to review an identified abnormality classifier category **444** and/or abnormality pattern category **446** in the inference data **1110** based on categorized performance data and/or qualification data of a corresponding user entry. The selected user can also correspond to a medical professional or other user employed at the originating entity and/or corresponding to the originating medical professional, indicated in the originating entity data **423**.

FIG. **6B** illustrates an embodiment of the medical scan diagnosing system **108** in conjunction with performing a remediation step **1140**. The medical scan diagnosing system **108** can monitor the performance of the set of medical scan inference functions **1105**, for example, based on evaluating inference accuracy data outputted by an inference data

evaluation function and/or based monitoring on the performance score data **630** in the medical scan analysis function database, and can determine whether or not if the corresponding medical scan inference function **1105** is performing properly. This can include, for example, determining if a remediation step **1140** is necessary for a medical scan inference function **1105**, for example, by comparing the performance score data **630** and/or inference accuracy data to remediation criteria data **652**. Determining if a remediation step **1140** is necessary can also be based on receiving an indication from the expert user or another user that remediation is necessary for one or more identified medical scan inference functions **1105** and/or for all of the medical scan inference functions **1105**.

In various embodiments, a remediation evaluation function is utilized to determine if a remediation step **1140** is necessary for medical scan inference function **1105**. The remediation evaluation function can include determining that remediation is necessary when recent accuracy data and/or efficiency data of a particular medical scan inference function **1105** is below the normal performance level of the particular inference function. The remediation evaluation function can include determining that remediation is necessary when recent or overall accuracy data and/or efficiency data of a particular medical scan inference function **1105** is below a recent or overall average for all or similar medical scan inference functions **1105**. The remediation evaluation function can include determining that remediation is necessary only after a threshold number of incorrect diagnoses are made. In various embodiments, multiple threshold number of incorrect diagnoses correspond to different diagnoses categories. For example, the threshold number of incorrect diagnoses for remediation can be higher for false negative diagnoses than false positive diagnoses. Similarly, categories corresponding to different diagnosis severities and/or rarities can have different thresholds, for example where a threshold number of more severe and/or more rare diagnoses that were inaccurate to necessitate remediation is lower than a threshold number of less severe and/or less rare diagnoses that were inaccurate.

The remediation step **1140** can include automatically updating an identified medical inference function **1105**. This can include automatically retraining identified medical inference function **1105** on the same training set or on a new training set that includes new data, data with higher corresponding confidence scores, or data selected based on new training set criteria. The identified medical inference function **1105** can also be updated and/or changed based on the review data received from the client device. For example, the medical scan and expert feedback data can be added to the training set of the medical scan inference function **1105**, and the medical scan inference function **1105** can be retrained on the updated training set. Alternatively or in addition, the expert user can identify additional parameters and/or rules in the expert feedback data based on the errors made by the inference function in generating the inference data **1110** for the medical scan, and these parameters and/or rules can be applied to update the medical scan inference function, for example, by updating the model type data **622** and/or model parameter data **623**.

The remediation step **1140** can also include determining to split a scan category **1120** into two or more subcategories. Thus, two or more new medical scan inference functions **1105** can be created, where each new medical scan inference functions **1105** is trained on a corresponding training set that is a subset of the original training set and/or includes new medical scan data corresponding to the subcategory. This

31

can allow medical scan inference functions **1105** to become more specialized and/or allow functions to utilize characteristics and/or discrepancies specific to the subcategory when generating inference data **1110**. Similarly, a new scan category **1120** that was not previously represented by any of the medical scan inference functions **1105** can be added in the remediation step, and a new medical scan inference functions **1105** can be trained on a new set of medical scan data that corresponds to the new scan category **1120**. Splitting a scan category and/or adding a scan category can be determined automatically by the medical scan diagnosing system **108** when performing the remediation step **1140**, for example, based on performance score data **630**. This can also be determined based on receiving instructions to split a category and/or add a new scan category from the expert user or other user of the system.

After a medical scan inference function **1105** is updated or created for the first time, the remediation step **1140** can further undergo a commissioning test, which can include rigorous testing of the medical scan inference function **1105** on a testing set, for example, based on the training parameters **620**. For example, the commissioning test can be passed when the medical scan inference function **1105** generates a threshold number of correct inference data **1110** and/or the test can be passed if an overall or average discrepancy level between the inference data and the test data is below a set error threshold. The commissioning test can also evaluate efficiency, where the medical scan inference function **1105** only passes the commissioning test if it performs at or exceeds a threshold efficiency level. If the medical scan inference function **1105** fails the commissioning test, the model type and/or model parameters can be modified automatically or based on user input, and the medical scan inference function can be retested, continuing this process until the medical scan inference function **1105** passes the commissioning test.

The remediation step **1140** can include decommissioning the medical scan inference function **1105**, for example, while the medical scan inference function is being retrained and/or is undergoing the commissioning test. Incoming scans to the medical scan diagnosing system **108** with a scan category **1120** corresponding to a decommissioned medical scan inference function **1105** can be sent directly to review by one or more users, for example, in conjunction with the medical scan annotator system **106**. These user-reviewed medical scans and corresponding annotations can be included in an updated training set used to train the decommissioned medical scan inference function **1105** as part of the remediation step **1140**. In some embodiments, previous versions of the plurality of medical scan image analysis functions can be stored in memory of the medical scan diagnosing system and/or can be determined based on the version data **640** of a medical scan inference function **1105**. A previous version of a medical scan inference function **1105**, such as most recent version or version with the highest performance score, can be utilized during the remediation step **1140** as an alternative to sending all medical scans to user review.

A medical scan inference function can also undergo the remediation step **1140** automatically in response to a hardware and/or software update on processing, memory, and/or other computing devices where the medical scan inference function **1105** is stored and/or performed. Different medical scan inference functions **1105** can be containerized on their own devices by utilizing a micro-service architecture, so hardware and/or software updates may only necessitate that one of the medical scan inference functions **1105** undergo

32

the remediation step **1140** while the others remain unaffected. A medical scan inference function **1105** can also undergo the remediation step **1140** automatically in response to normal system boot-up, and/or periodically in fixed intervals. For example, in response to a scheduled or automatically detected hardware and/or software update, change, or issue, one or more medical scan inference functions **1105** affected by this hardware or software can be taken out of commission until they each pass the commissioning test. Such criteria can be indicated in the remediation criteria data **652**.

The medical scan diagnosing system **108** can automatically manage usage data, subscription data, and/or billing data for the plurality of users corresponding to user usage of the system, for example, by utilizing, generating, and/or updating some or all of the subscription data of the user database. Users can pay for subscriptions to the system, which can include different subscription levels that can correspond to different costs. For example, a hospital can pay a monthly cost to automatically diagnose up to 100 medical scans per month. The hospital can choose to upgrade their subscription or pay per-scan costs for automatic diagnosing of additional scans received after the quota is reached and/or the medical scan diagnosing system **108** can automatically send medical scans received after the quota is reached to an expert user associated with the hospital. In various embodiments incentive programs can be used by the medical scan diagnosing system to encourage experts to review medical scans from different medical entities. For example, an expert can receive credit to their account and/or subscription upgrades for every medical scan reviewed, or after a threshold number of medical scans are reviewed. The incentive programs can include interactions by a user with other subsystems, for example, based on contributions made to medical scan entries via interaction with other subsystems.

FIG. 7A presents an embodiment of a medical scan image analysis system **112**. A training set of medical scans used to train one more medical scan image analysis functions can be received from one or more client devices via the network and/or can be retrieved from the medical scan database **342**, for example, based on training set data **621** corresponding to medical scan image analysis functions. Training set criteria, for example, identified in training parameters **620** of the medical scan image analysis function, can be utilized to automatically identify and select medical scans to be included in the training set from a plurality of available medical scans. The training set criteria can be automatically generated based on, for example, previously learned criteria, and/or training set criteria can be received via the network, for example, from an administrator of the medical scan image analysis system. The training set criteria can include a minimum training set size. The training set criteria can include data integrity requirements for medical scans in the training set such as requiring that the medical scan is assigned a truth flag **461**, requiring that performance score data for a hospital and/or medical professional associated with the medical scan compares favorably to a performance score threshold, requiring that the medical scan has been reviewed by at least a threshold number of medical professionals, requiring that the medical scan and/or a diagnosis corresponding to a patient file of the medical scan is older than a threshold elapsed time period, or based on other criteria intended to insure that the medical scans and associated data in the training set is reliable enough to be considered "truth" data. The training set criteria can include longitudinal requirements such the number of required sub-

sequent medical scans for the patient, multiple required types of additional scans for the patient, and/or other patient file requirements.

The training set criteria can include quota and/or proportion requirements for one or more medical scan classification data. For example, the training set criteria can include meeting quota and/or proportion requirements for one or more scan types and/or human body location of scans, meeting quota or proportion requirements for a number of normal medical scans and a number of medical scans with identified abnormalities, meeting quota and/or proportion requirements for a number of medical scans with abnormalities in certain locations and/or a number of medical scans with abnormalities that meet certain size, type, or other characteristics, meeting quota and/or proportion data for a number of medical scans with certain diagnosis or certain corresponding medical codes, and/or meeting other identified quota and/or proportion data relating to metadata, patient data, or other data associated with the medical scans.

In some embodiments, multiple training sets are created to generate corresponding medical scan image analysis functions, for example, corresponding to some or all of the set of medical scan inference functions **1105**. Some or all training sets can be categorized based on some or all of the scan classifier data **420** as described in conjunction with the medical scan diagnosing system **108**, where medical scans are included in a training set based on their scan classifier data **420** matching the scan category of the training set. In some embodiments, the input quality assurance function **1106** or another input check step can be performed on medical scans selected for each training set to confirm that their corresponding scan classifier data **420** is correct. In some embodiments, the input quality assurance function can correspond to its own medical scan image analysis function, trained by the medical scan image analysis system, where the input quality assurance function utilizes high level computer vision technology to determine a scan category **1120** and/or to confirm the scan classifier data **420** already assigned to the medical scan.

In some embodiments, the training set will be used to create a single neural network model, or other model corresponding to model type data **622** and/or model parameter data **623** of the medical scan image analysis function that can be trained on some or all of the medical scan classification data described above and/or other metadata, patient data, or other data associated with the medical scans. In other embodiments, a plurality of training sets will be created to generate a plurality of corresponding neural network models, where the multiple training sets are divided based on some or all of the medical scan classification data described above and/or other metadata, patient data, or other data associated with the medical scans. Each of the plurality of neural network models can be generated based on the same or different learning algorithm that utilizes the same or different features of the medical scans in the corresponding one of the plurality of training sets. The medical scan classifications selected to segregate the medical scans into multiple training sets can be received via the network, for example based on input to an administrator client device from an administrator. The medical scan classifications selected to segregate the medical scans can be automatically determined by the medical scan image analysis system, for example, where an unsupervised clustering algorithm is applied to the original training set to determine appropriate medical scan classifications based on the output of the unsupervised clustering algorithm.

In embodiments where the medical scan image analysis system is used in conjunction with the medical scan diagnosing system, each of the medical scan image analysis functions associated with each neural network model can correspond to one of the plurality of neural network models generated by the medical scan image analysis system. For example, each of the plurality of neural network models can be trained on a training set classified on scan type, scan human body location, hospital or other originating entity data, machine model data, machine calibration data, contrast agent data, geographic region data, and/or other scan classifying data as discussed in conjunction with the medical scan diagnosing system. In embodiments where the training set classifiers are learned, the medical scan diagnosing system can determine which of the medical scan image analysis functions should be applied based on the learned classifying criteria used to segregate the original training set.

A computer vision-based learning algorithm used to create each neural network model can include selecting a three-dimensional subregion **1310** for each medical scan in the training set. This three-dimensional subregion **1310** can correspond to a region that is "sampled" from the entire scan that may represent a small fraction of the entire scan. Recall that a medical scan can include a plurality of ordered cross-sectional image slices. Selecting a three-dimensional subregion **1310** can be accomplished by selecting a proper image slice subset **1320** of the plurality of cross-sectional image slices from each of the plurality of medical scans, and by further selecting a two-dimensional subregion **1330** from each of the selected subset of cross-sectional image slices of the each of the medical scans. In some embodiments, the selected image slices can include one or more non-consecutive image slices and thus a plurality of disconnected three-dimensional subregions will be created. In other embodiments, the selected proper subset of the plurality of image slices correspond to a set of consecutive image slices, as to ensure that a single, connected three-dimensional subregion is selected. In some embodiments, entire scans of the training set are used to train the neural network model. In such embodiment, as used herein, the three-dimensional subregion **1310** can refer to all of the medical scan image data **410** of a medical scan.

In some embodiments, a density windowing step can be applied to the full scan or the selected three-dimensional subregion. The density windowing step can include utilizing a selected upper density value cut off and/or a selected lower density value cut off, and masking pixels with higher values than the upper density value cut off and/or masking pixels with lower values than the lower density value cut off. The upper density value cut off and/or a selected lower density value cut off can be determined based on based on the range and/or distribution of density values included in the region that includes the abnormality, and/or based on the range and/or distribution of density values associated with the abnormality itself, based on user input to a subsystem, based on display parameter data associated with the medical scan or associated with medical scans of the same type, and/or can be learned in the training step. In some embodiments, a non-linear density windowing function can be applied to alter the pixel density values, for example, to stretch or compress contrast. In some embodiments, this density windowing step can be performed as a data augmenting step, to create additional training data for a medical scan in accordance with different density windows.

Having determined the subregion training set **1315** of three-dimensional subregions **1310** corresponding to the set of full medical scans in the training set, the medical scan

image analysis system can complete a training step **1352** by performing a learning algorithm on the plurality of three-dimensional subregions to generate model parameter data **1355** of a corresponding learning model. The learning model can include one or more of a neural network, an artificial neural network, a convolutional neural network, a Bayesian model, a support vector machine model, a cluster analysis model, or other supervised or unsupervised learning model. The model parameter data **1355** can be generated by performing the learning algorithm **1350**, and the model parameter data **1355** can be utilized to determine the corresponding medical scan image analysis functions. For example, some or all of the model parameter data **1355** can be mapped to the medical scan analysis function in the model parameter data **623** or can otherwise define the medical scan analysis function.

The training step **1352** can include creating feature vectors for each three-dimensional subregion of the training set for use by the learning algorithm **1350** to generate the model parameter data **1355**. The feature vectors can include the pixel data of the three-dimensional subregions such as density values and/or grayscale values of each pixel based on a determined density window. The feature vectors can also include other features as additional input features or desired output features, such as known abnormality data such as location and/or classification data, patient history data such as risk factor data or previous medical scans, diagnosis data, responsible medical entity data, scan machinery model or calibration data, contrast agent data, medical code data, annotation data that can include raw or processed natural language text data, scan type and/or anatomical region data, or other data associated with the image, such as some or all data of a medical scan entry **352**. Features can be selected based on administrator instructions received via the network and/or can be determined based on determining a feature set that reduces error in classifying error, for example, by performing a cross-validation step on multiple models created using different feature sets. The feature vector can be split into an input feature vector and output feature vector. The input feature vector can include data that will be available in subsequent medical scan input, which can include for example, the three-dimensional subregion pixel data and/or patient history data. The output feature vector can include data that will be inferred in in subsequent medical scan input and can include single output value, such as a binary value indicating whether or not the medical scan includes an abnormality or a value corresponding to one of a plurality of medical codes corresponding to the image. The output feature vector can also include multiple values which can include abnormality location and/or classification data, diagnosis data, or other output. The output feature vector can also include a determined upper density value cut off and/or lower density value cut off, for example, characterizing which pixel values were relevant to detecting and/or classifying an abnormality. Features included in the output feature vector can be selected to include features that are known in the training set, but may not be known in subsequent medical scans such as triaged scans to be diagnosed by the medical scan diagnosing system, and/or scans to be labeled by the medical scan report labeling system. The set of features in the input feature vector and output feature vector, as well as the importance of different features where each feature is assigned a corresponding weight, can also be designated in the model parameter data **1355**.

Consider a medical scan image analysis function that utilizes a neural network. The neural network can include a

plurality of layers, where each layer includes a plurality of neural nodes. Each node in one layer can have a connection to some or all nodes in the next layer, where each connection is defined by a weight value. Thus, the model parameter data **1355** can include a weight vector that includes weight values for every connection in the network. Alternatively or in addition, the model parameter data **1355** can include any vector or set of parameters associated with the neural network model, which can include an upper density value cut off and/or lower density value cut off used to mask some of the pixel data of an incoming image, kernel values, filter parameters, bias parameters, and/or parameters characterizing one or more of a plurality of convolution functions of the neural network model. The medical scan image analysis function can be utilized to produce the output vector as a function of the input feature vector and the model parameter data **1355** that characterizes the neural network model. In particular, the medical scan image analysis function can include performing a forward propagation step plurality of neural network layers to produce an inferred output vector based on the weight vector or other model parameter data **1355**. Thus, the learning algorithm **1350** utilized in conjunction with a neural network model can include determining the model parameter data **1355** corresponding to the neural network model, for example, by populating the weight vector with optimal weights that best reduce output error.

In particular, determining the model parameter data **1355** can include utilizing a backpropagation strategy. The forward propagation algorithm can be performed on at least one input feature vector corresponding to at least one medical scan in the training set to propagate the at least one input feature vector through the plurality of neural network layers based on initial and/or default model parameter data **1355**, such as an initial weight vector of initial weight values set by an administrator or chosen at random. The at least one output vector generated by performing the forward propagation algorithm on the at least one input feature vector can be compared to the corresponding at least one known output feature vector to determine an output error. Determining the output error can include, for example, computing a vector distance such as the Euclidian distance, or squared Euclidian distance, between the produced output vector and the known output vector, and/or determining an average output error such as an average Euclidian distance or squared Euclidian distance if multiple input feature vectors were employed. Next, gradient descent can be performed to determine an updated weight vector based on the output error or average output error. This gradient descent step can include computing partial derivatives for the error with respect to each weight, or other parameter in the model parameter data **1355**, at each layer starting with the output layer. Chain rule can be utilized to iteratively compute the gradient with respect to each weight or parameter at each previous layer until all weight's gradients are computed. Next updated weights, or other parameters in the model parameter data **1355**, are generated by updating each weight based on its corresponding calculated gradient. This process can be repeated on at least one input feature vector, which can include the same or different at least one feature vector used in the previous iteration, based on the updated weight vector and/or other updated parameters in the model parameter data **1355** to create a new updated weight vector and/or other new updated parameters in the model parameter data **1355**. This process can continue to repeat until the output error converges, the output error is within a certain error threshold, or another criterion is reached to determine the most recently

updated weight vector and/or other model parameter data **1355** is optimal or otherwise determined for selection.

Having determined the medical scan neural network and its final other model parameter data **1355**, an inference step **1354** can be performed on new medical scans to produce inference data **1370**, such as inferred output vectors, as shown in FIG. 7B. The inference step can include performing the forward propagation algorithm to propagate an input feature vector through a plurality of neural network layers based on the final model parameter data **1355**, such as the weight values of the final weight vector, to produce the inference data. This inference step **1354** can correspond to performing the medical scan image analysis function, as defined by the final model parameter data **1355**, on new medical scans to generate the inference data **1370**, for example, in conjunction with the medical scan diagnosing system **108** to generate inferred diagnosis data or other selected output data for triaged medical scans based on its corresponding the input feature vector.

The inference step **1354** can include applying the density windowing step to new medical scans. Density window cut off values and/or a non-linear density windowing function that are learned can be automatically applied when performing the inference step. For example, if the training step **1352** was used to determine optimal upper density value cut off and/or lower density value cut off values to designate an optimal density window, the inference step **1354** can include masking pixels of incoming scans that fall outside of this determined density window before applying the forward propagation algorithm. As another example, if learned parameters of one or more convolutional functions correspond to the optimal upper density value cut off and/or lower density value cut off values, the density windowing step is inherently applied when the forward propagation algorithm is performed on the new medical scans.

In some embodiments where a medical scan analysis function is defined by model parameter data **1355** corresponding to a neural network model, the neural network model can be a fully convolutional neural network. In such embodiments, only convolution functions are performed to propagate the input feature vector through the layers of the neural network in the forward propagation algorithm. This enables the medical scan image analysis functions to process input feature vectors of any size. For example, as discussed herein, the pixel data corresponding to the three-dimensional subregions is utilized input to the forward propagation algorithm when the training step **1352** is employed to populate the weight vector and/or other model parameter data **1355**. However, when performing the forward propagation algorithm in the inference step **1354**, the pixel data of full medical scans can be utilized as input, allowing the entire scan to be processed to detect and/or classify abnormalities, or otherwise generate the inference data **1370**. This may be a preferred embodiment over other embodiments where new scans must also be sampled by selecting a three-dimensional subregions and/or other embodiments where the inference step requires "piecing together" inference data **1370** corresponding to multiple three-dimensional subregions processed separately.

The inferred output vector of the inference data **1370** can include a plurality of abnormality probabilities mapped to a pixel location of each of a plurality of cross-sectional image slices of the new medical scan. For example, the inferred output vector can indicate a set of probability matrices **1371**, where each matrix in the set corresponds to one of the plurality of image slices of the medical scan, where each matrix is a size corresponding to the number of pixels in

each image slice, where each cell of each matrix corresponds to a pixel of the corresponding image slice, whose value is the abnormality probability of the corresponding pixel.

A detection step **1372** can include determining if an abnormality is present in the medical scan based on the plurality of abnormality probabilities. Determining if an abnormality is present can include, for example, determining that a cluster of pixels in the same region of the medical scan correspond to high abnormality probabilities, for example, where a threshold proportion of abnormality probabilities must meet or exceed a threshold abnormality probability, where an average abnormality probability of pixels in the region must meet or exceed a threshold abnormality probability, where the region that includes the cluster of pixels must be at least a certain size, etc. Determining if an abnormality is present can also include calculating a confidence score based on the abnormality probabilities and/or other data corresponding to the medical scan such as patient history data. The location of the detected abnormality can be determined in the detection step **1372** based on the location of the pixels with the high abnormality probabilities. The detection step can further include determining an abnormality region **1373**, such as a two-dimensional subregion on one or more image slices that includes some or all of the abnormality. The abnormality region **1373** determined in the detection step **1372** can be mapped to the medical scan to populate some or all of the abnormality location data **443** for use by one or more other subsystems **101** and/or client devices **120**. Furthermore, determining whether or not an abnormality exists in the detection step **1372** can be used to populate some or all of the diagnosis data **440** of the medical scan, for example, to indicate that the scan is normal or contains an abnormality in the diagnosis data **440**.

An abnormality classification step **1374** can be performed on a medical scan in response to determining an abnormality is present. Classification data **1375** corresponding to one or more classification categories such as abnormality size, volume, pre-post contract, doubling time, calcification, components, smoothness, texture, diagnosis data, one or more medical codes, a malignancy rating such as a Lung-RADS score, or other classifying data as described herein can be determined based on the detected abnormality. The classification data **1375** generated by the abnormality classification step **1374** can be mapped to the medical scan to populate some or all of the abnormality classification data **445** of the corresponding abnormality classifier categories **444** and/or abnormality pattern categories **446** and/or to determine one or more medical codes **447** of the medical scan. The abnormality classification step **1374** can include performing an abnormality classification function on the full medical scan, or the abnormality region **1373** determined in the detection step **1372**. The abnormality classification function can be based on another model trained on abnormality data such as a support vector machine model, another neural network model, or any supervised classification model trained on medical scans, or portions of medical scans, that include known abnormality classifying data to generate inference data for some or all of the classification categories. For example, the abnormality classification function can include another medical scan analysis function. Classification data **1375** in each of a plurality of classification categories can also be assigned their own calculated confidence score, which can also be generated by utilizing the abnormality classification function. Output to the abnormality classification function can also include at least one identified similar medical scan and/or at least one identified similar cropped image, for example, based on the training data. The

abnormality classification step can also be included in the inference step **1354**, where the inferred output vector or other inference data **1370** of the medical scan image analysis function includes the classification data **1375**.

The abnormality classification function can be trained on full medical scans and/or one or more cropped or full selected image slices from medical scans that contain an abnormality. For example, the abnormality classification function can be trained on a set of two-dimensional cropped slices that include abnormalities. The selected image slices and/or the cropped region in each selected image slice for each scan in the training set can be automatically selected based upon the known location of the abnormality. Input to the abnormality classification function can include the full medical scan, one or more selected full image slices, and/or one or more selected image slices cropped based on a selected region. Thus, the abnormality classification step can include automatically selecting one or more image slices that include the detected abnormality. The slice selection can include selecting the center slice in a set of consecutive slices that are determined to include the abnormality or selecting a slice that has the largest cross-section of the abnormality, or selecting one or more slices based on other criteria. The abnormality classification step can also include automatically generating one or more cropped two-dimensional images corresponding to the one or more of the selected image slices based on an automatically selected region that includes the abnormality.

Input to the abnormality classification function can also include other data associated with the medical scan, including patient history, risk factors, or other metadata. The abnormality classification step can also include determining some or all of the characteristics based on data of the medical scan itself. For example, the abnormality size and volume can be determined based on a number of pixels determined to be part of the detected abnormality. Other classifiers such as abnormality texture and/or smoothness can be determined by performing one or more other preprocessing functions on the image specifically designed to characterize such features. Such preprocessed characteristics can be included in the input to the abnormality classification function to the more difficult task of assigning a medical code or generating other diagnosis data. The training data can also be preprocessed to include such preprocessed features.

A similar scan identification step **1376** can also be performed on a medical scan with a detected abnormality and/or can be performed on the abnormality region **1373** determined in the detection step **1372**. The similar scan identification step **1376** can include generating similar abnormality data **1377**, for example, by identifying one or more similar medical scans or one or more similar cropped two-dimensional images from a database of medical scans and/or database of cropped two-dimensional images. Similar medical scans and/or cropped images can include medical scans or cropped images that are visually similar, medical scans or cropped images that have known abnormalities in a similar location to an inferred abnormality location of the given medical scan, medical scans that have known abnormalities with similar characteristics to inferred characteristics of an abnormality in the given scan, medical scans with similar patient history and/or similar risk factors, or some combination of these factors and/or other known and/or inferred factors. The similar abnormality data **1377** can be mapped to the medical scan to populate some or all of its corresponding similar scan data **480** for use by one or more other subsystems **101** and/or client devices **120**.

The similar scans identification step **1376** can include performing a scan similarity algorithm, which can include generating a feature vector for the given medical scan and for medical scans in the set of medical scans, where the feature vector can be generated based on quantitative and/or category based visual features, inferred features, abnormality location and/or characteristics such as the predetermined size and/or volume, patient history and/or risk factor features, or other known or inferred features. A medical scan similarity analysis function can be applied to the feature vector of the given medical scan and one or more feature vectors of medical scans in the set. The medical scan similarity analysis function can include computing a similarity distance such as the Euclidian distance between the feature vectors, and assigning the similarity distance to the corresponding medical scan in the set. Similar medical scans can be identified based on determining one or more medical scans in the set with a smallest computed similarity distance, based on ranking medical scans in the set based on the computed similarity distances and identifying a designated number of top ranked medical scans, and/or based on determining if a similarity distance between the given medical scan and a medical scan in the set is smaller than a similarity threshold. Similar medical scans can also be identified based on determining medical scans in a database that mapped to a medical code that matches the medical code of the medical scan, or mapped to other matching classifying data. A set of identified similar medical scans can also be filtered based on other inputted or automatically generated criteria, where for example only medical scans with reliable diagnosis data or rich patient reports, medical scans with corresponding with longitudinal data in the patient file such as multiple subsequent scans taken at later dates, medical scans with patient data that corresponds to risk factors of the given patient, or other identified criteria, where only a subset of scans that compare favorably to the criteria are selected from the set and/or only a highest ranked single scan or subset of scans are selected from the set, where the ranking is automatically computed based on the criteria. Filtering the similar scans in this fashion can include calculating, or can be based on previously calculated, one or more scores as discussed herein. For example, the ranking can be based on a longitudinal quality score, such as the longitudinal quality score **434**, which can be calculated for an identified medical scan based on a number of subsequent and/or previous scans for the patient. Alternatively or in addition, the ranking can be based on a confidence score associated with diagnosis data of the scan, such as confidence score data **460**, based on performance score data associated with a user or medical entity associated with the scan, based on an amount of patient history data or data in the medical scan entry **352**, or other quality factors. The identified similar medical scans can be filtered based on ranking the scans based on their quality score and/or based on comparing their quality score to a quality score threshold. In some embodiments, a longitudinal threshold must be reached, and only scans that compare favorably to the longitudinal threshold will be selected. For example, only scans with at least three scans on file for the patient and final biopsy data will be included.

In some embodiments, the similarity algorithm can be utilized in addition to or instead of the trained abnormality classification function to determine some or all of the inferred classification data **1375** of the medical scan, based on the classification data such as abnormality classification data **445** or other diagnosis data **440** mapped to one or more of the identified similar scans. In other embodiments, the similarity algorithm is merely used to identify similar scans

for review by medical professionals to aid in review, diagnosis, and/or generating medical reports for the medical image.

A display parameter step **1378** can be performed based on the detection and/or classification of the abnormality. The display parameter step can include generating display parameter data **1379**, which can include parameters that can be used by an interactive interface to best display each abnormality. The same or different display parameters can be generated for each abnormality. The display parameter data generated in the display parameter step **1378** can be mapped to the medical scan to populate some or all of its corresponding display parameter data **470** for use by one or more other subsystems **101** and/or client devices **120**.

Performing the display parameter step **1378** can include selecting one or more image slices that include the abnormality by determining the one or more image slices that include the abnormality and/or determining one or more image slices that has a most optimal two-dimensional view of the abnormality, for example by selecting the center slice in a set of consecutive slices that are determined to include the abnormality, selecting a slice that has the largest cross-section of the abnormality, selecting a slice that includes a two-dimensional image of the abnormality that is most similar to a selected most similar two-dimensional-image, selecting the slice that was used as input to the abnormality classification step and/or similar scan identification step, or based on other criteria. This can also include automatically cropping one or more selected image slices based on an identified region that includes the abnormality. This can also select an ideal Hounsfield window that best displays the abnormality. This can also include selecting other display parameters based on data generated by the medical scan interface evaluating system and based on the medical scan.

FIGS. **8A-8F** illustrate embodiments of a medical picture archive integration system **2600**. The medical picture archive integration system **2600** can provide integration support for a medical picture archive system **2620**, such as a PACS that stores medical scans. The medical picture archive integration system **2600** can utilize model parameters received from a central server system **2640** via a network **2630** to perform an inference function on de-identified medical scans of medical scans received from the medical picture archive system **2620**. The annotation data produced by performing the inference function can be transmitted back to the medical picture archive system. Furthermore, the annotation data and/or de-identified medical scans can be sent to the central server system **2640**, and the central server system can train on this information to produce new and/or updated model parameters for transmission back to the medical picture archive integration system **2600** for use on subsequently received medical scans.

In various embodiments, medical picture archive integration system **2600** includes a de-identification system that includes a first memory designated for protected health information (PHI), operable to perform a de-identification function on a DICOM image, received from a medical picture archive system, to identify at least one patient identifier and generate a de-identified medical scan that does not include the at least one patient identifier. The medical picture archive integration system further includes a de-identified image storage system that stores the de-identified medical scan in a second memory that is separate from the first memory, and an annotating system, operable to utilize model parameters received from a central server to perform an inference function on the de-identified medical scan,

retrieved from the second memory to generate annotation data for transmission to the medical picture archive system as an annotated DICOM file.

The first memory and the second memory can be implemented by utilizing separate storage systems: the first memory can be implemented by a first storage system designated for PHI storage, and the second memory can be implemented by a second storage system designated for storage of de-identified data. The first storage system can be protected from access by the annotating system, while the second storage system can be accessible by the annotating system. The medical picture archive integration system **2600** can be operable to perform the de-identification function on data in first storage system to generate de-identified data. The de-identified data can then be stored in the second storage system for access by the annotating system. The first and second storage systems can be physically separate, each utilizing at least one of their own, separate memory devices. Alternatively, the first and second storage systems can be virtually separate, where data is stored in separate virtual memory locations on the same set of memory devices. Firewalls, virtual machines, and/or other protected containerization can be utilized to enforce the separation of data in each storage system, to protect the first storage system from access by the annotating system and/or from other unauthorized access, and/or to ensure that only data of the first storage system that has been properly de-identified through application of the de-identification function can be stored in the second storage system.

As shown in FIG. **8A**, the medical picture archive system **2620** can receive image data from a plurality of modality machines **2622**, such as CT machines, MRI machines, x-ray machines, and/or other medical imaging machines that produce medical scans. The medical picture archive system **2620** can store this image data in a DICOM image format and/or can store the image data in a plurality of medical scan entries **352** as described in conjunction with some or all of the attributes described in conjunction with FIGS. **4A** and **4B**. While "DICOM image" will be used herein to refer to medical scans stored by the medical picture archive system **2620**, the medical picture archive integration system **2600** can provide integration support for medical picture archive systems **2620** that store medical scans in other formats.

The medical picture archive integration system **2600** can include a receiver **2602** and a transmitter **2604**, operable to transmit and receive data from the medical picture archive system **2620**, respectively. For example, the receiver **2602** and transmitter **2604** can be configured to receive and transmit data, respectively, in accordance with a DICOM communication protocol and/or another communication protocol recognized by the medical image archive system **2620**. The receiver can receive DICOM images from the medical picture archive system **2620**. The transmitter **2604** can send annotated DICOM files to the medical picture archive system **2620**.

DICOM images received via receiver **2602** can be sent directly to a de-identification system **2608**. The de-identification system **2608** can be operable to perform a de-identification function on the first DICOM image to identify at least one patient identifier in the DICOM image, and to generate a de-identified medical scan that does not include the identified at least one patient identifier. As used herein, a patient identifier can include any patient identifying data in the image data, header, and/or metadata of a medical scan, such as a patient ID number or other unique patient identifier, an accession number, a service-object pair (SOP) instance unique identifier (UID) field, scan date and/or time

that can be used to determine the identity of the patient that was scanned at that date and/or time, and/or other private data corresponding to the patient, doctor, or hospital. In some embodiments, the de-identified medical scan is still in a DICOM image format. For example, a duplicate DICOM image that does not include the patient identifiers can be generated, and/or the original DICOM image can be altered such that the patient identifiers of the new DICOM image are masked, obfuscated, removed, replaced with a custom fiducial, and/or otherwise anonymized. In other embodiments, the de-identified medical scan is formatted in accordance with a different image format and/or different data format that does not include the identifying information. In some embodiments, other private information, for example, associated with a particular doctor or other medical professional, can be identified and anonymized as well.

Some patient identifying information can be included in a DICOM header of the DICOM image, for example, in designated fields for patient identifiers. These corresponding fields can be anonymized within the corresponding DICOM header field. Other patient identifying information can be included in the image itself, such as in medical scan image data **410**. For example, the image data can include a patient name or other identifier that was handwritten on a hard copy of the image before the image was digitized. As another example, a hospital administered armband or other visual patient information in the vicinity of the patient may have been captured in the image itself. A computer vision model can detect the presence of these identifiers for anonymization, for example, where a new DICOM image includes a fiducial image that covers the identifying portion of the original DICOM image. In some embodiments, patient information identified in the DICOM header can be utilized to detect corresponding patient information in the image itself. For example, a patient name extracted from the DICOM header before anonymization can be used to search for the patient name in the image and/or to detect a location of the image that includes the patient name. In some embodiments, the de-identification system **2608** is implemented by the de-identification system discussed in conjunction with FIGS. **10A**, **10B** and **11**, and/or utilizes functions and/or operations discussed in conjunction with FIGS. **10A**, **10B** and **11**.

The de-identified medical scan can be stored in de-identified image storage system **2610** and the annotating system **2612** can access the de-identified medical scan from the de-identified image storage system **2610** for processing. The de-identified storage system can archive a plurality of de-identified DICOM images and/or can serve as temporary storage for the de-identified medical scan until processing of the de-identified medical scan by the annotating system **2612** is complete. The annotating system **2612** can generate annotation data by performing an inference function on the de-identified medical scan, utilizing the model parameters received from the central server system **2640**. The annotation data can correspond to some or all of the diagnosis data **440** as discussed in conjunction with FIGS. **4A** and **4B**. In some embodiments, the annotating system **2612** can utilize the model parameters to perform inference step **1354**, the detection step **1372**, the abnormality classification step **1374**, the similar scan identification step **1376**, and/or the display parameter step **1378** of the medical scan image analysis system **112**, as discussed in conjunction with FIG. **7B**, on de-identified medical scans received from the medical picture archive system **2620**.

In some embodiments, model parameters for a plurality of inference functions can be received from the central server

system **2640**, for example, where each inference function corresponds to one of a set of different scan categories. Each scan category can correspond to a unique combination of one or a plurality of scan modalities, one of a plurality of anatomical regions, and/or other scan classifier data **420**. For example, a first inference function can be trained on and intended for de-identified medical scans corresponding chest CT scans, and a second inference function can be trained on and intended for de-identified medical scans corresponding to head MM scans. The annotating system can select one of the set of inference functions based on determining the scan category of the DICOM image, indicated in the de-identified medical scan, and selecting the inference function that corresponds to the determined scan category.

To ensure that scans received from the medical picture archive system **2620** match the set of scan categories for which the annotating system is operable to perform a corresponding inference function, the transmitter can transmit requests, such as DICOM queries, indicating image type parameters such as parameters corresponding to scan classifier data **420**, for example indicating one or more scan modalities, one or more anatomical regions, and/or other parameters. For example, the request can indicate that all incoming scans that match the set of scan categories corresponding to a set of inference functions the annotating system **2612** for which the annotating system has obtained model parameters from the central server system **2640** and is operable to perform.

Once the annotation data is generated by performing the selected inference function, the annotating system **2612** can generate an annotated DICOM file for transmission to the medical image archive system **2620** for storage. The annotated DICOM file can include some or all of the fields of the diagnosis data **440** and/or abnormality annotation data **442** of FIGS. **4A** and **4B**. The annotated DICOM file can include scan overlay data, providing location data of an identified abnormality and/or display data that can be used in conjunction with the original DICOM image to indicate the abnormality visually in the DICOM image and/or otherwise visually present the annotation data, for example, for use with the medical scan assisted review system **102**. For example, a DICOM presentation state file can be generated to indicate the location of an abnormality identified in the de-identified medical scan. The DICOM presentation state file can include an identifier of the original DICOM image, for example, in metadata of the DICOM presentation state file, to link the annotation data to the original DICOM image. In other embodiments, a full, duplicate DICOM image is generated that includes the annotation data with an identifier linking this duplicate annotated DICOM image to the original DICOM image.

The identifier linking the annotated DICOM file to the original DICOM image can be extracted from the original DICOM file by the de-identification system **2608**, thus enabling the medical picture archive system **2620** to link the annotated DICOM file to the original DICOM image in its storage. For example, the de-identified medical scan can include an identifier that links the de-identified medical scan to the original DICOM file, but does not link the de-identified medical scan to a patient identifier or other private data.

In some embodiments, generating the annotated DICOM file includes altering one or more fields of the original DICOM header. For example, standardized header formatting function parameters can be received from the central server system and can be utilized by the annotating system to alter the original DICOM header to match a standardized

DICOM header format. The standardized header formatting function can be trained in a similar fashion to other medical scan analysis functions discussed herein and/or can be characterized by some or all fields of a medical scan analysis function entry 356. The annotating system can perform the standardized header formatting function on a de-identified medical scan to generate a new, standardized DICOM header for the medical scan to be sent back to the medical picture archive system 2620 in the annotated DICOM file and/or to replace the header of the original DICOM file. The standardized header formatting function can be run in addition to other inference functions utilized to generate annotation data. In other embodiments, the medical picture archive integration system 2600 is implemented primarily for header standardization for medical scans stored by the medical picture archive system 2620. In such embodiments, only the standardized header formatting function is performed on the de-identified data to generate a modified DICOM header for the original DICOM image, but the de-identified medical scan is not annotated.

In some embodiments of header standardization, the annotation system can store a set of acceptable, standardized entries for some or all of the DICOM header fields, and can select one of the set of acceptable, standardized entries in populating one or more fields of the new DICOM header for the annotated DICOM file. For example, each of the set of scan categories determined by the annotating system can correspond to a standardized entry of one or more fields of the DICOM header. The new DICOM header can thus be populated based on the determined scan category.

In some embodiments, each of the set of standardized entries can be mapped to a set of related, non-standardized entries, such as entries in a different order, commonly misspelled entries, or other similar entries that do not follow a standardized format. For example, one of the set of acceptable, standardized entries for a field corresponding to a scan category can include "Chest CT", which can be mapped to a set of similar, non-standardized entries which can include "CT chest", "computerized topography CT", and/or other entries that are not standardized. In such embodiments, the annotating system can determine the original DICOM header is one of the similar non-standardized entries, and can select the mapped, standardized entry as the entry for the modified DICOM header. In other embodiments, the image data itself and/or other header data can be utilized by the annotation system to determine a standardized field. For example, an input quality assurance function 1106 can be trained by the central server system and sent to the annotating system to determine one or more appropriate scan classifier fields, or one or more other DICOM header fields, based on the image data or other data of the de-identified medical scan. One or more standardized labels can be assigned to corresponding fields of the modified DICOM header based on the one or more fields determined by the input quality assurance function.

In some embodiments, the DICOM header is modified based on the annotation data generated in performing the inference function. In particular, a DICOM priority header field can be generated and/or modified automatically based on the severity and/or time-sensitivity of the abnormalities detected in performing the inference function. For example, a DICOM priority header field can be changed from a low priority to a high priority in response to annotation data indicating a brain bleed in the de-identified medical scan of a DICOM image corresponding to a head CT scan, and a new DICOM header that includes the high priority DICOM priority header field can be sent back to the medical picture

archive system 2620 to replace or otherwise be mapped to the original DICOM image of the head CT scan.

In various embodiments, the medical picture archive system 2620 is disconnected from network 2630, for example, to comply with requirements regarding Protected Health Information (PHI), such as patient identifiers and other private patient information included in the DICOM images and/or otherwise stored by the medical picture archive system 2620. The medical picture archive integration system 2600 can enable processing of DICOM images while still protecting private patient information by first de-identifying DICOM data by utilizing de-identification system 2608. The de-identification system 2608 can utilize designated processors and memory of the medical picture archive integration system, for example, designated for PHI. The de-identification system 2608 can be decoupled from the network 2630 to prevent the DICOM images that still include patient identifiers from being accessed via the network 2630. For example, as shown in FIG. 8A, the de-identification system 2608 is not connected to network interface 2606. Furthermore, only the de-identification system 2608 has access to the original DICOM files received from the medical picture archive system 2620 via receiver 2602. The de-identified image storage system 2610 and annotating system 2612, as they are connected to network 2630 via network interface 2606, only store and have access to the de-identified medical scan produced by the de-identification system 2608.

This containerization that separates the de-identification system 2608 from the de-identified image storage system 2610 and the annotating system 2612 is further illustrated in FIG. 8B, which presents an embodiment of the medical picture archive integration system 2600. The de-identification system 2608 can include its own designated memory 2654 and processing system 2652, connected to receiver 2602 via bus 2659. For example, this memory 2654 and processing system 2652 can be designated for PHI, and can adhere to requirements for handling PHI. The memory 2654 can store executable instructions that, when executed by the processing system 2652, enable the de-identification system to perform the de-identification function on DICOM images received via receiver 2602 of the de-identification system. The incoming DICOM images can be temporarily stored in memory 2654 for processing, and patient identifiers detected in performing the de-identification function can be temporarily stored in memory 2654 to undergo anonymization. Interface 2655 can transmit the de-identified medical scan to interface 2661 for use by the de-identified image storage system 2610 and the annotating system 2612. Interface 2655 can be protected from transmitting original DICOM files and can be designated for transmission of de-identified medical scan only.

Bus 2669 connects interface 2661, as well as transmitter 2604 and network interface 2606, to the de-identified image storage system 2610 and the annotating system 2612. The de-identified image storage system 2610 and annotating system 2612 can utilize separate processors and memory, or can utilize shared processors and/or memory. For example, the de-identified image storage system 2610 can serve as temporary memory of the annotating system 2612 as de-identified images are received and processed to generate annotation data.

As depicted in FIG. 8B, the de-identified image storage system 2610 can include memory 2674 that can temporarily store incoming de-identified medical scans as it undergoes processing by the annotating system 2612 and/or can archive a plurality of de-identified medical scans corresponding to a

plurality of DICOM images received by the medical picture archive integration system **2600**. The annotating system **2612** can include a memory **2684** that stores executable instructions that, when executed by processing system **2682**, cause the annotating system **2612** to perform a first inference function on de-identified medical scan to generate annotation data by utilizing the model parameters received via interface **2606**, and to generate an annotated DICOM file based on the annotation data for transmission via transmitter **2604**. The model parameters can be stored in memory **2684**, and can include model parameters for a plurality of inference functions, for example, corresponding to a set of different scan categories.

The medical picture archive integration system can be an onsite system, installed at a first geographic site, such as a hospital or other medical entity that is affiliated with the medical picture archive system **2620**. The hospital or other medical entity can further be responsible for the PHI of the de-identification system, for example, where the memory **2654** and processing system **2652** are owned by, maintained by, and/or otherwise affiliated with the hospital or other medical entity. The central server system **2640** can be located at a second, separate geographic site that is not affiliated with the hospital or other medical entity and/or at a separate geographic site that is not affiliated with the medical picture archive system **2620**. The central server system **2640** can be a server configured to be outside the network firewall and/or out outside the physical security of the hospital or other medical entity or otherwise not covered by the particular administrative, physical and technical safeguards of the hospital or other medical entity.

FIG. **8C** further illustrates how model parameters can be updated over time to improve existing inference functions and/or to add new inference functions, for example corresponding to new scan categories. In particular, the some or all of the de-identified medical scans generated by the de-identification system **2608** can be transmitted back to the central server system, and the central server system **2640** can train on this data to improve existing models by producing updated model parameters of an existing inference function and/or to generate new models, for example, corresponding to new scan categories, by producing new model parameters for new inference functions. For example, the central server system **2640** can produce updated and/or new model parameters by performing the training step **1352** of the medical scan image analysis system **112**, as discussed in conjunction with FIG. **7A**, on a plurality of de-identified medical scans received from the medical picture archive integration system **2600**.

The image type parameters can be determined by the central server system to dictate characteristics of the set of de-identified medical scans to be received to train and/or retrain the model. For example, the image type parameters can correspond to one or more scan categories, can indicate scan classifier data **420**, can indicate one or more scan modalities, one or more anatomical regions, a date range, and/or other parameters. The image type parameters can be determined by the central server system based on training parameters **620** determined for the corresponding inference function to be trained, and/or based on characteristics of a new and/or existing scan category corresponding to the inference function to be trained. The image type parameters can be sent to the medical picture archive integration system **2600**, and a request such as a DICOM query can be sent to the medical picture archive system **2620**, via transmitter **2604**, that indicates the image type parameters. For example, the processing system **2682** can be utilized to generate the

DICOM query based on the image type parameters received from the central server system **2640**. The medical picture archive system can automatically transmit one or more DICOM images to the medical picture archive integration system in response to determining that the one or more DICOM images compares favorably to the image type parameters. The DICOM images received in response can be de-identified by the de-identification system **2608**. In some embodiments, the de-identified medical scans can be transmitted directly to the central server system **2640**, for example, without generating annotation data.

The central server system can generate the new and/or updated model parameters by training on the received set of de-identified medical scans, and can transmit the new and/or updated model parameters to the de-identified storage system. If the model parameters correspond to a new inference function for a new scan category, the medical picture archive integration system **2600** can generate a request, such as a DICOM query, for transmission to the medical picture archive system indicating that incoming scans corresponding to image type parameters corresponding to the new scan category be sent to the medical picture archive integration system. The annotating system can update the set of inference functions to include the new inference function, and the annotating system can select the new inference function from the set of inference functions for subsequently generated de-identified medical scans by the de-identification system by determining each of these de-identified medical scans indicate the corresponding DICOM image corresponds to the new scan category. The new model parameters can be utilized to perform the new inference function on each of these de-identified medical scans to generate corresponding annotation data, and an annotated DICOM file corresponding to each of these de-identified medical scans can be generated for transmission to the medical picture archive system via the transmitter.

In some embodiments, the central server system **2640** receives a plurality of de-identified medical scans from a plurality of medical picture archive integration system **2600**, for example, each installed at a plurality of different hospitals or other medical entities, via the network **2630**. The central server system can generate training sets by integrating de-identified medical scans from some or all of the plurality of medical picture archive integration systems **2600** to train one or more inference functions and generate model parameters. The plurality of medical picture archive integration systems **2600** can utilize the same set of inference functions or different sets of inference functions. In some embodiments, the set of inference functions utilized by the each of the plurality of medical picture archive systems **2620** are trained on different sets of training data. For example, the different sets of training data can correspond to the set of de-identified medical scans received from the corresponding medical picture archive integration system **2600**.

In some embodiments, the medical scan diagnosing system **108** can be utilized to implement the annotating system **2612**, where the corresponding subsystem processing device **235** and subsystem memory device **245** of the medical scan diagnosing system **108** are utilized to implement the processing system **2682** and the memory **2684**, respectively. Rather than receiving the medical scans via the network **150** as discussed in conjunction with FIG. **6A**, the medical scan diagnosing system **108** can perform a selected medical scan inference function **1105** on an incoming de-identified medical scan generated by the de-identification system **2608** and/or retrieved from the de-identified image storage system **2610**. Memory **2684** can store the set of medical scan

inference functions **1105**, each corresponding to a scan category **1120**, where the inference function is selected from the set based on determining the scan category of the de-identified medical scan and selecting the corresponding inference function. The processing system **2682** can perform the selected inference function **1105** to generate the inference data **1110**, which can be further utilized by the annotating system **2612** to generate the annotated DICOM file for transmission back to the medical picture archive system **2620**. New medical scan inference functions **1105** can be added to the set when corresponding model parameters are received from the central server system. The remediation step **1140** can be performed locally by the annotating system **2612** and/or can be performed by the central server system **2640** by utilizing one or more de-identified medical scans and corresponding annotation data sent to the central server system **2640**. Updated model parameters can be generated by the central server system **2640** and sent to the medical picture archive integration system **2600** as a result of performing the remediation step **1140**.

The central server system **2640** can be implemented by utilizing one or more of the medical scan subsystems **101**, such as the medical scan image analysis system **112** and/or the medical scan diagnosing system **108**, to produce model parameters for one or more inference functions. The central server system can store or otherwise communicate with a medical scan database **342** that includes the de-identified medical scans and/or annotation data received from one or more medical picture archive integration systems **2600**. Some or all entries of the medical scan database **342** can be utilized to as training data to produce model parameters for one or more inference functions. These entries of the medical scan database **342** can be utilized by other subsystems **101** as discussed herein. For example, other subsystems **101** can utilize the central server system **2640** to fetch medical scans and/or corresponding annotation data that meet specified criteria. The central server system **2640** can query the medical picture archive integration system **2600** based on this criteria, and can receive de-identified medical scans and/or annotation data in response. This can be sent to the requesting subsystem **101** directly and/or can be added to the medical scan database **342** or another database of the database storage system **140** for access by the requesting subsystem **101**.

Alternatively or in addition, the central server system **2640** can store or otherwise communicate with a user database **344** storing user profile entries corresponding to each of a plurality of medical entities that each utilize a corresponding one of a plurality of medical picture archive integration systems **2600**. For example, basic user data corresponding to the medical entity can be stored as basic user data, a number of scans or other consumption information indicating usage of one or more inference functions by corresponding medical picture archive integration system can be stored as consumption usage data, and/or a number of scans or other contribution information indicating de-identified scans sent to the central server system as training data can be stored as contribution usage data. The user profile entry can also include inference function data, for example, with a list of model parameters or function identifiers, such as medical scan analysis function identifiers **357**, of inference functions currently utilized by the corresponding medical picture archive integration system **2600**. These entries of the user database **344** can be utilized by other subsystems **101** as discussed herein.

Alternatively or in addition, the central server system **2640** can store or otherwise communicate with a medical

scan analysis function database **346** to store model parameters, training data, or other information for one or more inference functions as medical scan analysis function entries **356**. In some embodiments, model parameter data **623** can indicate the model parameters and function classifier data **610** can indicate the scan category of inference function entries. In some embodiments, the medical scan analysis function entry **356** can further include usage identifying information indicating a medical picture archive integration system identifier, medical entity identifier, and/or otherwise indicating which medical archive integration systems and/or medical entities have received the corresponding model parameters to utilize the inference function corresponding to the medical scan analysis function entry **356**. These entries of the medical scan analysis function database **346** can be utilized by other subsystems **101** as discussed herein.

In some embodiments, the de-identification function is a medical scan analysis function, for example, with a corresponding medical scan analysis function entry **356** in the medical scan analysis function database **346**. In some embodiments, the de-identification function is trained by the central server system **2640**. For example, the central server system **2640** can send de-identification function parameters to the medical picture archive integration system **2600** for use by the de-identification system **2608**. In embodiments with a plurality of medical picture archive integration systems **2600**, each of the plurality of medical picture archive integration systems **2600** can utilize the same or different de-identification functions. In some embodiments, the de-identification function utilized by the each of the plurality of medical picture archive integration systems **2600** are trained on different sets of training data. For example, the different sets of training data can correspond to each different set of de-identified medical scans received from each corresponding medical picture archive integration system **2600**.

In some embodiments, as illustrated in FIGS. **8D-8F**, the medical picture archive integration system **2600** can further communicate with a report database **2625**, such as a Radiology Information System (RIS), that includes a plurality of medical reports corresponding to the DICOM images stored by the medical picture archive system **2620**.

As shown in FIG. **8D**, the medical picture archive integration system **2600** can further include a receiver **2603** that receives report data, corresponding to the DICOM image, from report database **2625**. The report database **2625** can be affiliated with the medical picture archive system **2620** and can store report data corresponding to DICOM images stored in the medical picture archive system. The report data of report database **2625** can include PHI, and the report database **2625** can thus be disconnected from network **2630**.

The report data can include natural language text, for example, generated by a radiologist that reviewed the corresponding DICOM image. The report data can be used to generate the de-identified medical scan, for example, where the de-identification system **2608** performs a natural language analysis function on the report data to identify patient identifying text in the report data. The de-identification system **2608** can utilize this patient identifying text to detect matching patient identifiers in the DICOM image to identify the patient identifiers of the DICOM image and generate the de-identified medical scan. In some embodiments, the report data can be de-identified by obfuscating, hashing, removing, replacing with a fiducial, or otherwise anonymizing the identified patient identifying text to generate de-identified report data.

The de-identified report data can be utilized by the annotating system **2612**, for example, in conjunction with

51

the DICOM image, to generate the annotation data. For example, the annotating system **2612** can perform a natural language analysis function on the de-identified natural language text of the report data to generate some or all of the annotation data. In some embodiments, the de-identified report data is sent to the central server system, for example, to be used as training data for inference functions, for natural language analysis functions, for other medical scan analysis functions, and/or for use by at least one other subsystem **101**. For example, other subsystems **101** can utilize the central server system **2640** to fetch medical reports that correspond to particular medical scans or otherwise meet specified criteria. The central server system **2640** can query the medical picture archive integration system **2600** based on this criteria, and can receive de-identified medical reports in response. This can be sent to the requesting subsystem **101** directly, can be added to the medical scan database **342**, a de-identified report database, or another database of the database storage system **140** for access by the requesting subsystem **101**.

In some embodiments the medical picture archive integration system **2600** can query the report database **2625** for the report data corresponding to a received DICOM image by utilizing a common identifier extracted from the DICOM image.

In some embodiments, the report data can correspond to a plurality of DICOM images. For example, the report data can include natural language text describing a plurality of medical scans of a patient that can include multiple sequences, multiple modalities, and/or multiple medical scans taken over time. In such embodiments, the patient identifying text and/or annotation data detected in the report data can also be applied to de-identify and/or generate annotation data for the plurality of DICOM images it describes. In such embodiments, the medical picture archive integration system **2600** can query the medical picture archive system **2620** for one or more additional DICOM images corresponding to the report data, and de-identified data and annotation data for these additional DICOM images can be generated accordingly by utilizing the report data.

In some embodiments, as shown in FIG. 8E, the medical picture archive system **2620** communicates with the report database **2625**. The medical picture archive system **2620** can request the report data corresponding to the DICOM image from the report database **2625**, and can transmit the report data to the medical picture archive integration system **2600** via a DICOM communication protocol for receipt via receiver **2602**. The medical picture archive system **2620** can query the report database **2625** for the report data, utilizing a common identifier extracted from the corresponding DICOM image, in response to determining to send the corresponding DICOM image to the medical picture archive integration system **2600**.

FIG. 8F presents an embodiment where report data is generated by the annotating system **2612** and is transmitted, via a transmitter **2605**, to the report database **2625**, for example via a DICOM communication protocol or other protocol recognized by the report database **2625**. In other embodiments, the report data is instead transmitted via transmitter **2604** to the medical picture archive system **2620**, and the medical picture archive system **2620** transmits the report data to the report database **2625**.

The report data can be generated by the annotating system **2612** as output of performing the inference function on the de-identified medical scan. The report data can include natural language text data **448** generated automatically based on other diagnosis data **440** such as abnormality

52

annotation data **442** determined by performing the inference function, for example, by utilizing a medical scan natural language generating function trained by the medical scan natural language analysis system **114**. The report data can be generated instead of, or in addition to, the annotated DICOM file.

FIG. 9 presents a flowchart illustrating a method for execution by a medical picture archive integration system **2600** that includes a first memory and a second memory that store executional instructions that, when executed by at least one first processor and at least one second processor, respectfully, cause the medical picture archive integration system to perform the steps below. In various embodiments, the first memory and at least one first processor are implemented by utilizing, respectfully, the memory **2654** and processing system **2652** of FIG. 8B. In various embodiments, the second memory is implemented by utilizing the memory **2674** and/or the memory **2684** of FIG. 8B. In various embodiments, the at least one second processor is implemented by utilizing the processing system **2682** of FIG. 8B.

Step **2702** includes receiving, from a medical picture archive system via a receiver, a first DICOM image for storage in the first memory, designated for PHI, where the first DICOM image includes at least one patient identifier. Step **2704** includes performing, via at least one first processor coupled to the first memory and designated for PHI, a de-identification function on the first DICOM image to identify the at least one patient identifier and generate a first de-identified medical scan that does not include the at least one patient identifier.

Step **2706** includes storing the first de-identified medical scan in a second memory that is separate from the first memory. Step **2708** includes receiving, via a network interface communicating with a network that does not include the medical picture archive system, first model parameters from a central server.

Step **2710** includes retrieving the first de-identified medical scan from the second memory. Step **2712** includes utilizing the first model parameters to perform a first inference function on the first de-identified medical scan to generate first annotation data via at least one second processor that is different from the at least one first processor. Step **2714** includes generating, via the at least one second processor, a first annotated DICOM file for transmission to the medical picture archive system via a transmitter, where the first annotated DICOM file includes the first annotation data and further includes an identifier that indicates the first DICOM image. In various embodiments, the first annotated DICOM file is a DICOM presentation state file.

In various embodiments, the second memory further includes operational instructions that, when executed by the at least one second processor, further cause the medical picture archive integration system to retrieve a second de-identified medical scan from the de-identified image storage system, where the second de-identified medical scan was generated by the at least one first processor by performing the de-identification function on a second DICOM image received from the medical picture archive system. The updated model parameters are utilized to perform the first inference function on the second de-identified medical scan to generate second annotation data. A second annotated DICOM file is generated for transmission to the medical picture archive system via the transmitter, where the second annotated DICOM file includes the second annotation data and further includes an identifier that indicates the second DICOM image.

In various embodiments, the second memory stores a plurality of de-identified medical scans generated by the at least one first processor by performing the de-identification function on a corresponding plurality of DICOM images received from the medical picture archive system via the receiver. The plurality of de-identified medical scans is transmitted to the central server via the network interface, and the central server generates the first model parameters by performing a training function on training data that includes the plurality of de-identified medical scans.

In various embodiments, the central server generates the first model parameters by performing a training function on training data that includes a plurality of de-identified medical scans received from a plurality of medical picture archive integration systems via the network. Each of the plurality of medical picture archive integration systems communicates bidirectionally with a corresponding one of a plurality of medical picture archive systems, and the plurality of de-identified medical scans corresponds to a plurality of DICOM images stored by the plurality of medical picture archive integration systems.

In various embodiments, the first de-identified medical scan indicates a scan category of the first DICOM image. The second memory further stores operational instructions that, when executed by the at least one second processor, further cause the medical picture archive integration system to select the first inference function from a set of inference functions based on the scan category. The set of inference functions corresponds to a set of unique scan categories that includes the scan category. In various embodiments, each unique scan category of the set of unique scan categories is characterized by one of a plurality of modalities and one of a plurality of anatomical regions.

In various embodiments, the first memory further stores operational instructions that, when executed by the at least one first processor, further cause the medical picture archive integration system to receive a plurality of DICOM image data from the medical picture archive system via the receiver for storage in the first memory in response to a query transmitted to the medical picture archive system via the transmitter. The query is generated by the medical picture archive integration system in response to a request indicating a new scan category received from the central server via the network. The new scan category is not included in the set of unique scan categories, and the plurality of DICOM image data corresponds to the new scan category. The de-identification function is performed on the plurality of DICOM image data to generate a plurality of de-identified medical scans for transmission to the central server via the network.

The second memory further stores operational instructions that, when executed by the at least one second processor, further cause the medical picture archive integration system to receive second model parameters from the central server via the network for a new inference function corresponding to the new scan category. The set of inference functions is updated to include the new inference function. The second de-identified medical scan is retrieved from the first memory, where the second de-identified medical scan was generated by the at least one first processor by performing the de-identification function on a second DICOM image received from the medical picture archive system. The new inference function is selected from the set of inference functions by determining the second de-identified medical scan indicates the second DICOM image corresponds to the new scan category. The second model parameters are utilized to perform the new inference function on the second

de-identified medical scan to generate second annotation data. A second annotated DICOM file is generated for transmission to the medical picture archive system via the transmitter, where the second annotated DICOM file includes the second annotation data and further includes an identifier that indicates the second DICOM image.

In various embodiments, the medical picture archive integration system generates parameter data for transmission to the medical picture archive system that indicates the set of unique scan categories. The medical picture archive system automatically transmits the first DICOM image to the medical picture archive integration system in response to determining that the first DICOM image compares favorably to one of the set of unique scan categories.

In various embodiments, the second memory further stores operational instructions that, when executed by the at least one second processor, cause the medical picture archive integration system to generate a natural language report data is based on the first annotation data and to transmit, via a second transmitter, the natural language report data to a report database associated with the medical picture archive integration system, where the natural language report data includes an identifier corresponding to the first DICOM image.

In various embodiments, the first memory further stores operational instructions that, when executed by the at least one first processor, cause the medical picture archive integration system to receive, via a second receiver, a natural language report corresponding to the first DICOM image from the report database. A set of patient identifying text included in the natural language report are identified. Performing the de-identification function on the first DICOM image includes searching the first DICOM image for the set of patient identifying text to identify the at least one patient identifier.

In various embodiments, the first memory is managed by a medical entity associated with the medical picture archive system. The medical picture archive integration system is located at a first geographic site corresponding to the medical entity, and the central server is located at a second geographic site. In various embodiments, the first memory is decoupled from the network to prevent the first DICOM image that includes the at least one patient identifier from being communicated via the network. In various embodiments, the medical picture archive system is a Picture Archive and Communication System (PACS) server, and the first DICOM image is received in response to a query sent to the medical picture archive system by the transmitter in accordance with a DICOM communication protocol.

FIG. 10A presents an embodiment of a de-identification system **2800**. The de-identification system **2800** can be utilized to implement the de-identification system **2608** of FIGS. **8A-8F**. In some embodiments, the de-identification system **2800** can be utilized by other subsystems to de-identify image data, medical report data, private fields of medical scan entries **352** such as patient identifier data **431**, and/or other private fields stored in databases of the database memory device **340**.

The de-identification system can be operable to receive, from at least one first entity, a medical scan and a medical report corresponding to the medical scan. A set of patient identifiers can be identified in a subset of fields of a header of the medical scan. A header anonymization function can be performed on each of the set of patient identifiers to generate a corresponding set of anonymized fields. A de-identified

medical scan can be generated by replacing the subset of fields of the header of the medical scan with the corresponding set of anonymized fields.

A subset of patient identifiers of the set of patient identifiers can be identified in the medical report by searching text of the medical report for the set of patient identifiers. A text anonymization function can be performed on the subset of patient identifiers to generate corresponding anonymized placeholder text for each of the subset of patient identifiers. A de-identified medical report can be generated by replacing each of the subset of patient identifiers with the corresponding anonymized placeholder text. The de-identified medical scan and the de-identified medical report can be transmitted to a second entity via a network.

As shown in FIG. 10A, the de-identification system **2800** can include at least one receiver **2802** operable to receive medical scans, such as medical scans in a DICOM image format. The at least one receiver **2802** is further operable to receive medical reports, such as report data **449** or other reports containing natural language text diagnosing, describing, or otherwise associated the medical scans received by the de-identification system. The medical scans and report data can be received from the same or different entity, and can be received by the same or different receiver **2802** in accordance with the same or different communication protocol. For example, the medical scans can be received from the medical picture archive system **2620** of FIGS. **8A-8F** and the report data can be received from the report database **2625** of FIGS. **8D-8F**. In such embodiments, the receiver **2802** can be utilized to implement the receiver **2602** of FIG. **8B**.

The de-identification system **2800** can further include a processing system **2804** that includes at least one processor, and a memory **2806**. The memory **2806** can store operational instructions that, when executed by the processing system, cause the de-identification system to perform at least one patient identifier detection function on the received medical scan and/or the medical report to identify a set of patient identifiers in the medical scan and/or the medical report. The operational instructions, when executed by the processing system, can further cause the de-identification system to perform an anonymization function on the medical scan and/or the medical report to generate a de-identified medical scan and/or a de-identified medical report that do not include the set of patient identifiers found in performing the at least one patient identifier detection function. Generating the de-identified medical scan can include generating a de-identified header and generating de-identified image data, where the de-identified medical scan includes both the de-identified header and the de-identified image data. The memory **2806** can be isolated from Internet connectivity, and can be designated for PHI.

The de-identification system **2800** can further include at least one transmitter **2808**, operable to transmit the de-identified medical scan and de-identified medical report. The de-identified medical scan and de-identified medical report can be transmitted back to the same entity from which they were received, respectively, and/or can be transmitted to a separate entity. For example, the at least one transmitter can transmit the de-identified medical scan to the de-identified image storage system **2610** of FIGS. **8A-8F** and/or can transmit the de-identified medical scan to central server system **2640** via network **2630** of FIGS. **8A-8F**. In such embodiments, the transmitter **2808** can be utilized to implement the interface **2655** of FIG. **8B**. The receiver **2802**, processing system **2804**, memory **2806**, and/or transmitter **2808** can be connected via bus **2810**.

Some or all of the at least one patient identifier detection function and/or at least one anonymization function as discussed herein can be trained and/or implemented by one or subsystems **101** in the same fashion as other medical scan analysis functions discussed herein, can be stored in medical scan analysis function database **346** of FIG. **3**, and/or can otherwise be characterized by some or all fields of a medical scan analysis function entry **356** of FIG. **5**.

The de-identification system **2800** can perform separate patient identifier detection functions on the header of a medical report and/or medical scan, on the text data of the medical report, and/or on the image data of the medical scan, such as text extracted from the image data of the medical scan. Performance of each of these functions generates an output of its own set of identified patient identifiers. Combining these sets of patient identifiers yields a blacklist term set. A second pass of the header of a medical report and/or medical scan, on the text data of the medical report, and/or on the image data of the medical scan that utilizes this blacklist term set can catch any terms that were missed by the respective patient identifier detection function, and thus, the outputs of these multiple identification processes can support each other. For example, some of the data in the headers will be in a structured form and can thus be easier to reliably identify. This can be exploited and used to further anonymize these identifiers when they appear in free text header fields, report data, and/or in the image data of the medical scan. Meanwhile, unstructured text in free text header fields, report data, and/or image data of the medical scan likely includes pertinent clinical information to be preserved in the anonymization process, for example, so it can be leveraged by at least one subsystems **101** and/or so it can be leveraged in training at least one medical scan analysis function.

At least one first patient identifier detection function can include extracting the data in a subset of fields of a DICOM header, or another header or other metadata of the medical scan and/or medical report with a known type that corresponds to patient identifying data. For example, this patient identifying subset of fields can include a name field, a patient ID number field or other unique patient identifier field, a date field, a time field, an age field, an accession number field, SOP instance UID, and/or other fields that could be utilized to identify the patient and/or contain private information. A non-identifying subset of fields of the header can include hospital identifiers, machine model identifiers, and/or some or all fields of medical scan entry **352** that do not correspond to patient identifying data. The patient identifying subset of fields and the non-identifying subset of fields can be mutually exclusive and collectively exhaustive with respect to the header. The at least one patient identifier function can include generating a first set of patient identifiers by ignoring the non-identifying subset of fields and extracting the entries of the patient identifying subset of fields only. This first set of patient identifiers can be anonymized to generate a de-identified header as discussed herein.

In some embodiments, at least one second patient identifier detection function can be performed on the report data of the medical report. The at least one second patient identifier detection function can include identifying patient identifying text in the report data by performing a natural language analysis function, for example, trained by the medical scan natural language analysis system **114**. For example, the at least one second patient identifier detection function can leverage the known structure of the medical report and/or context of the medical report. A second set of

patient identifiers corresponding to the patient identifying text can be determined, and the second set of patient identifiers can be anonymized to generate a de-identified medical report. In some embodiments, a de-identified medical report includes clinical information, for example, because the portion of the original medical report that includes the clinical information was deemed to be free of patient identifying text and/or because the portion of the original medical report that includes the clinical information was determined to include pertinent information to be preserved.

In some embodiments, the medical report includes image data corresponding to freehand or typed text. For example the medical report can correspond to a digitized scan of original freehand text written by a radiologist or other medical professional. In such embodiments, the patient identifier detection function can first extract the text from the freehand text in the image data to generate text data before the at least one second patient identifier detection function is performed on the text of the medical report to generate the second set of patient identifiers.

In some embodiments, the at least one second patient identifier detection function can similarly be utilized to identify patient identifying text in free text fields and/or unstructured text fields of a DICOM header and/or other metadata of the medical scan and/or medical report data by performing a natural language analysis function, for example, trained by the medical scan natural language analysis system **114**. A third set of patient identifiers corresponding to this patient identifying text of the free text and/or unstructured header fields can be determined, and the third set of patient identifiers can be anonymized to generate de-identified free text header field and/or unstructured header fields. In some embodiments, a de-identified free text header field and/or unstructured header field includes clinical information, for example, because the portion of the original corresponding header field that includes the clinical information was deemed to be free of patient identifying text and/or because the portion of the original corresponding header field that includes the clinical information was determined to include pertinent information to be preserved.

Patient identifiers can also be included in the image data of the medical scan itself. For example, freehand text corresponding to a patient name written on a hard copy of the medical scan before digitizing can be included in the image data, as discussed in conjunction with FIG. **10B**. Other patient identifiers, such as information included on a patient wristband or other identifying information located on or within the vicinity of the patient may have been captured when the medical scan was taken, and can thus be included in the image. At least one third patient identifier detection function can include extracting text from the image data and/or detecting non-text identifiers in the image data by performing a medical scan image analysis function, for example, trained by the medical scan image analysis system **112**. For example, detected text that corresponds to an image location known to include patient identifiers, detected text that corresponds to a format of a patient identifier, and/or detected text or other image data determined to correspond to a patient identifier can be identified. The at least one third patient identifier detection function can further include identifying patient identifying text in the text extracted from the image data by performing the at least one second patient identifier detection function and/or by performing a natural language analysis function. A fourth set of patient identifiers corresponding to patient identifying text or other patient identifiers detected in the image data of the medical scan can

be determined, and the fourth set of patient identifiers can be anonymized in the image data to generate de-identified image data of the medical scan as described herein. In particular, the fourth set of patient identifiers can be detected in a set of regions of image data of the medical scan, and the set of regions of the image data can be anonymized.

In some embodiments, only a subset of the patient identifier detection functions described herein are performed to generate respective sets of patient identifiers for anonymization. In some embodiments, additional patient identifier detection functions can be performed on the medical scan and/or medical report to determine additional respective sets of patient identifiers for anonymization. The sets of patient identifiers outputted by performing each patient identifier detection function can have a null or non-null intersection. The sets of patient identifiers outputted by performing each patient identifier function can have null or non-null set differences.

Cases where the sets of patient identifiers have non-null set differences can indicate that a patient identifier detected by one function may have been missed by another function. The combined set of patient identifiers, for example, generated as the union of the sets of sets of patient identifiers outputted by performing each patient identifier function, can be used to build a blacklist term set, for example, stored in memory **2806**. The blacklist term set can designate the final set of terms to be anonymized. A second pass of header data, medical scans, medical reports, and/or any free text extracted from the header data, the medical scan, and/or the medical report can be performed by utilizing the blacklist term set to flag terms for anonymization that were not caught in performing the respective at least one patient identifier detection function. For example, performing the second pass can include identifying at least one patient identifier of the blacklist term set in the header, medical report, and/or image data of the medical scan. This can include by searching corresponding extracted text of the header, medical report, and/or image data for terms included in blacklist term set and/or by determining if each term in the extracted text is included in the blacklist term set.

In some embodiments, at least one patient identifier is not detected until the second pass is performed. Consider an example where a free text field of a DICOM header included a patient name that was not detected in performing a respective patient identifier detection function on the free text field of the DICOM header. However, the patient name was successfully identified in the text of the medical report in performing a patient identifier detection function on the medical report. This patient name is added to the blacklist term list, and is detected in a second pass of the free text field of the DICOM header. In response to detection in the second pass, the patient name of the free text field of the DICOM header can be anonymized accordingly to generate a de-identified free text field. Consider a further example where the patient name is included in the image data of the medical scan, but was not detected in performing a respective patient identifier detection function on the free text field of the DICOM header. In the second pass, this patient name can be detected in at least one region of image data of the medical scan by searching the image data for the blacklist term set.

In some embodiments, performing some or all of the patient identifier detection functions includes identifying a set of non-identifying terms, such as the non-identifying subset of fields of the header. In particular, the non-identifying terms can include terms identified as clinical information and/or other terms determined to be preserved. The combined set of non-identifying terms, for example, gener-

ated as the union of the sets of sets of non-identifying outputted by performing each patient identifier function, can be used to build a whitelist term set, for example, stored in memory **2806**. Performing the second pass can further include identifying at least one non-identifying term of the whitelist term set in the header, medical report, and/or image data of the medical scan, and determining not to anonymize, or to otherwise ignore, the non-identifying term.

In various embodiments, some or all terms of the whitelist term set can be removed from the blacklist term set. In particular, at least one term previously identified as a patient identifier in performing one or more patient identifier detection functions is determined to be ignored and not anonymized in response to determining the term is included in the whitelist term set. This can help ensure that clinically important information is not anonymized, and is thus preserved in the de-identified medical scan and de-identified medical report.

In some embodiments, the second pass can be performed after each of the patient identifier detection functions are performed. For example, performing the anonymization function can include performing this second pass by utilizing the blacklist term set to determine the final set of terms to be anonymized. New portions of text in header fields, not previously detected in generating the first set of patient identifiers or the third set of patient identifiers, can be flagged for anonymization by determining these new portions of text correspond to terms of the blacklist term set. New portions of text the medical report, not previously detected in generating in the second set of patient identifiers, can be flagged for anonymization by determining these new portions of text correspond to terms of the blacklist term set. New regions of the image data of the medical scan, not previously detected in generating the fourth set of patient identifiers, can be flagged for anonymization by determining these new portions of text correspond to terms of the blacklist term set.

In some embodiments, the blacklist term set is built as each patient identifier detection function is performed, and performance of subsequent patient identifier detection functions includes utilizing the current blacklist term set. For example, performing the second patient identifier detection function can include identifying a first subset of the blacklist term set in the medical report by searching the text of the medical report for the blacklist term set and/or by determining if each term in the text of the medical report is included in the blacklist term set. Performing the second patient identifier detection function can further include identifying at least one term in the medical report that is included in the whitelist term set, and determining to ignore the term in response. The first subset can be anonymized to generate the de-identified medical report as discussed herein. New patient identifiers not already found can be appended to the blacklist term set, and the updated blacklist term set can be applied to perform a second search of the header and/or image data of the medical scan, and at least one of the new patient identifiers can be identified in the header in the second search of the header and/or in the image data in a second search of the image data. These newly identified patient identifiers in the header and/or image data are anonymized in generating the de-identified medical scan.

As another example, a second subset of the blacklist term set can be detected in a set of regions of image data of the medical scan by performing the medical scan image analysis function on image data of the medical scan, where the image analysis function includes searching the image data for the set of patient identifiers. For example, the medical scan

image analysis function can include searching the image data for text, and the second subset can include detected text that matches one or more terms of the blacklist term set. In some embodiments, detected text that matches one or more terms of the whitelist term set can be ignored. The second subset can be anonymized to generate de-identified image data as discussed herein. New patient identifiers that are detected can be appended to the blacklist term set, and the updated blacklist term set can be applied to perform a second search of the header and/or metadata of the medical scan, and/or can be applied to perform a second search of the medical report. At least one of the new patient identifiers can be identified in the header as a result of performing the second search of the header and/or at least one of the new patient identifiers can be identified medical report as a result of performing the second search of the medical report. These newly identified patient identifiers can be anonymized in the header along with the originally identified blacklist term set in generating the de-identified header, and/or can be anonymized in the medical report along with the originally identified first subset in generating the de-identified medical report.

In some embodiments, the memory **2806** further stores a global blacklist, for example, that includes a vast set of known patient identifying terms. In some embodiments, the global blacklist is also utilized by at least one patient identifier detection function and/or in performing the second pass to determine patient identifying terms for anonymization. In some embodiments, the blacklist term set generated for a particular medical scan and corresponding medical report can be appended to the global blacklist for use in performing the second pass and/or in detecting patient identifiers in subsequently received medical scans and/or medical reports.

Alternatively or in addition, the memory **2806** can further store a global whitelist, for example, that includes a vast set of terms that can be ignored. In particular, the global whitelist can include clinical terms and/or other terms that are deemed beneficial to preserve that do not correspond to patient identifying information. In some embodiments, the global whitelist is utilized by at least one patient identifier detection function and/or in performing the second pass to determine terms to ignore in the header, image data, and/or medical report. In some embodiments, the whitelist term set generated for a particular medical scan and corresponding medical report can be appended to the global whitelist for use in performing the second pass and/or in ignoring terms in subsequently received medical scans and/or medical reports.

Alternatively or in addition, the memory **2806** can further store a global graylist, for example, that includes ambiguous terms that could be patient identifying terms in some contexts, but non-identifying terms in other contexts. For example, "Parkinson" could correspond to patient identifying data if part of a patient name such as "John Parkinson", but could correspond to non-patient identifying data meant to be ignored and preserved in the de-identified medical report and/or de-identified medical scan if part of a diagnosis term such as "Parkinson's disease." In some embodiments, the global graylist is also utilized in performing the second pass and/or in performing at least one patient identifier detection function to determine that a term is included in the graylist, and to further determine whether the term should be added to the blacklist term set for anonymization or whitelist term set to be ignored by leveraging context of accompanying text, by leveraging known data types of a header field from which the term was extracted, by leveraging known

61

structure of the term, by leveraging known data types of a location of the image data from which the term was extracted, and/or by leveraging other contextual information. In some embodiments, the graylist term set can be updated based on blacklist and/or whitelist term sets for a particular medical scan and corresponding medical report.

In some embodiments, the at least one anonymization function includes a fiducial replacement function. For example, some or all of the blacklist term set can be replaced with a corresponding, global fiducial in the header, report data, and/or image data. In some embodiments, the global fiducial can be selected from a set of global fiducials based on a type of the corresponding patient identifier. Each patient identifier detected in the header and/or medical report can be replaced with a corresponding one of the set of global text fiducials. Each patient identifiers detected in the image data can be replaced with a corresponding one of the set of global image fiducials. For example, one or more global image fiducials can overlay pixels of regions of the image data that include the identifying patient data, to obfuscate the identifying patient data in the de-identified image data.

The global text fiducials and/or global image fiducials can be recognizable by inference functions and/or training functions, for example, where the global text fiducials and global image fiducials are ignored when processed in a training step to train an inference function and/or are ignored in an inference step when processed by an inference function. Furthermore, the global text fiducials and/or global image fiducials can be recognizable by a human viewing the header, medical report, and/or image data. For example, a radiologist or other medical professional, upon viewing a header, medical report, and/or image data, can clearly identify the location of a patient identifier that was replaced by the fiducial and/or can identify the type of patient identifier that was replaced by the fiducial.

As an example, the name "John Smith" can be replaced in a header and/or medical report with the text "% PATIENT NAME %", where the text "% PATIENT NAME %" is a global fiducial for name types of the header and/or the text of medical reports. The training step and/or inference step of medical scan natural language analysis functions can recognize and ignore text that matches "% PATIENT NAME %" automatically.

FIG. 10B illustrates an example of anonymizing patient identifiers in image data of a medical scan. In this example, the name "John Smith" and the date "May 4, 2010" is detected as freehand text in the original image data of a medical scan. The regions of the image data that include the patient identifiers can each be replaced by global fiducial in the shape of a rectangular bar, or any other shape. As shown in FIG. 10B, a first region corresponding to the location of "John Smith" in the original image data is replaced by fiducial 2820 in the de-identified image data, and a second region corresponding to the location of "May 4, 2010" in the original image data is replaced by fiducial 2822 in the de-identified image data. The size, shape, and/or location of each global visual fiducial can be automatically determined based on the size, shape, and/or location of the region that includes the patient identifier to minimize the amount of the image data that is obfuscated, while still ensuring the entirety of the text is covered. While not depicted in FIG. 10B, the fiducial can be of a particular color, for example, where pixels of the particular color are automatically recognized by the training step and/or inference step of medical scan image analysis functions to indicate that the corresponding region be ignored, and/or where the particular color is not included in the original medical scan and/or is

62

known to not be included in any medical scans. The fiducial can include text recognizable to human inspection such as "% PATIENT NAME" and "% DATE" as depicted in FIG. 10B, and/or can include a QR code, logo, or other unique symbol recognizable to human inspection and/or automatically recognizable by the training step and/or inference step of medical scan image analysis functions to indicate that the corresponding region be ignored.

In some embodiments, other anonymization functions can be performed on different ones of the patient identifying subset of fields to generate the de-identified header, de-identified report data, and/or de-identified image data. For example, based on the type of identifying data of each field of the header, different types of header anonymization functions and/or text anonymization functions can be selected and utilized on the header fields, text of the report, and/or text extracted from the image data. A set of anonymization functions can include a shift function, for example, utilized to offset a date, time or other temporal data by a determined amount to preserve absolute time difference and/or to preserve relative order over multiple medical scans and/or medical reports of a single patient. FIG. 10B depicts an example where the shift function is performed on the date detected in the image data to generate fiducial 2822, where the determined amount is 10 years and 1 month. The determined amount can be determined by the de-identification system randomly and/or pseudo-randomly for each patient and/or for each medical scan and corresponding medical report, ensuring the original date cannot be recovered by utilizing a known offset. In various embodiments, other medical scans and/or medical reports are fetched for the same patient by utilizing a patient ID number or other unique patient identifier of the header. These medical scans and reports can be anonymized as well, where the dates and/or times detected in these medical scans and/or medical reports offset by the same determined amount, randomized or pseudo-randomized for particular patient ID number, for example, based on performing a hash function on the patient ID number.

The set of anonymization functions can include at least one hash function, for example utilized to hash a unique patient ID such as a patient ID number, accession number, and/or SOP instance UID of the header and/or text. In some embodiments, the hashed SOP instance UID, accession number, and/or patient ID number are prepended with a unique identifier, stored in a database of the memory 2806 and/or shared with the entities to which the de-identified medical scans and/or medical reports are transmitted, so that de-identified medical scans and their corresponding de-identified medical reports can be linked and retrieved retroactively. Similarly, longitudinal data can be preserved as multiple medical scans and/or medical reports of the same patient will be assigned the same hashed patient ID.

The set of anonymization functions can further include at least one manipulator function for some types of patient identifiers. Some values of header fields and/or report text that would normally not be considered private information can be considered identifying patient data if they correspond to an outlier value or other rare value that could then be utilized to identify the corresponding patient from a very small subset of possible options. For example, a patient age over 89 could be utilized to determine the identity of the patient, for example, if there are very few patients over the age of 89. To prevent such cases, in response to determining that a patient identifier corresponds to an outlier value and/or in response to determining that a patient identifier compares unfavorably to a normal-range threshold value, the patient

63

identifier can be capped at the normal-range threshold value or can otherwise be manipulated. For example, a normal-range threshold value corresponding to age can be set at 89, and generating a de-identified patient age can include capping patient ages that are higher than 89 at 89 and/or can include keeping the same value for patient ages that are less than or equal to 89.

In some embodiments, the de-identified header data is utilized to replace the corresponding first subset of patient identifiers detected in the medical report with text of the de-identified header fields. In other embodiments, a set of text anonymization functions includes a global text fiducial replacement function, shift function, a hash function, and/or manipulator functions that anonymize the corresponding types of patient identifiers in the medical report separately.

In some embodiments where the image data of a medical scan includes an anatomical region corresponding to a patient's head, the image data may include an identifying facial structure and/or facial features that could be utilized to determine the patient's identity. For example, a database of facial images, mapped to a corresponding plurality of people including the patient, could be searched and a facial recognition function could be utilized to identify the patient in the database. Thus, facial structure included in the image data can be considered patient identifying data.

To prevent this problem and maintain patient privacy, the de-identification system can further be implemented to perform facial obfuscation for facial structure detected in medical scans. At least one region of the image data that includes identifying facial structure can be determined by utilizing a medical image analysis function. For example, the medical image analysis function can include a facial detection function that determines the regions of the image data that include identifying facial structure based on searching the image data for pixels with a density value that corresponds to facial skin, facial bone structure, or other density of an anatomical mass type that corresponds to identifying facial structure, and the facial obfuscation function can be performed on the identified pixels. Alternatively or in addition, the facial detection function can determine the region based on identifying at least one shape in the image data that corresponds to a facial structure.

The image obfuscation function can include a facial structure obfuscation function performed on the medical scan to generate de-identified image data that does not include identifying facial structure. For example, the facial structure obfuscation function can mask, scramble, replace with a fiducial, or otherwise obfuscate the pixels of the region identified by the facial detection function. In some embodiments, the facial structure obfuscation function can perform a one-way function on the region that preserves abnormalities of the corresponding portions of the image, such as nose fractures or facial skin lesions, while still obfuscating the identifying facial structure such that the patient is not identifiable. For example, the pixels of the identifying facial structure can be altered such that they converge towards a fixed, generic facial structure. In some embodiments, a plurality of facial structure image data of a plurality of patients can be utilized to generate the generic facial structure, for example, corresponding to an average or other combination of the plurality of faces. For example, the pixels of the generic facial structure can be averaged with, superimposed upon, or otherwise combined with the pixels of the region of the image data identified by the facial detection function in generating the de-identified image data.

64

In some embodiments, a hash function can be performed on an average of the generic facial structure and the identified facial structure of the image data so that the generic facial structure cannot be utilized in conjunction with the resulting data of the de-identified image data to reproduce the original, identifying facial structure. In such embodiments, the hash function can alter the pixel values while still preserving abnormalities. In some embodiments, a plurality of random, generic facial structures can be generated by utilizing the plurality of facial structure image data, for example, where each if the plurality of facial structure image data are assigned a random or pseudo-random weight in an averaging function utilized to create the generic facial structure, where a new, random or pseudo-random set of weights are generated each time the facial structure obfuscation function is utilized to create a new, generic facial structure to be averaged with the identified facial structure in creating the de-identified image data to ensure the original identifying facial structure cannot be extracted from the resulting de-identified image data.

While facial obfuscation is described herein, similar techniques can be applied in a similar fashion to other anatomical regions that are determined to include patient identifiers and/or to other anatomical regions that can be utilized to extract patient identifying information if not anonymized.

In some embodiments, the at least one receiver **2802** is included in at least one transceiver, for example, enabling bidirectional communication between the medical picture archive system **2620** and/or the report database **2625**. In such embodiments, the de-identification system **2800** can generate queries to the medical picture archive system **2620** and/or the report database **2625** for particular medical scans and/or medical reports, respectively. In particular, if the medical scan and medical report are stored and/or managed by separate memories and/or separate entities, they may not be received at the same time. However, a linking identifier, such as DICOM identifiers in headers or metadata of the medical scan and/or medical report, such as accession number, patient ID number, SOP instance UID, or other linking identifier that maps the medical scan to the medical report can be utilized to fetch a medical report corresponding to a received medical scan and/or to fetch a medical scan corresponding to a received medical report via a query sent utilizing the at least one transceiver. For example, in response to receiving the medical scan from the medical picture archive system **2620**, the de-identification system can extract a linking identifier from a DICOM header of the medical scan, and can query the report database **2625** for the corresponding medical report by indicating the linking identifier in the query. Conversely, in response to receiving the medical report from the report database **2625**, the de-identification system can extract the linking identifier from a header, metadata, and/or text body of the medical report, and can query the medical picture archive system **2620** for the corresponding medical scan by indicating the linking identifier in the query. In some embodiments, a mapping of de-identified medical scans to original medical scans, and/or a mapping of de-identified medical reports to original medical reports can be stored in memory **2806**. In some embodiments, linking identifiers such as patient ID numbers can be utilized to fetch additional medical scans, additional medical reports, or other longitudinal data corresponding to the same patient.

FIG. **11** presents a flowchart illustrating a method for execution by a de-identification system **2800** that stores

executorial instructions that, when executed by at least one processor, cause the de-identification to perform the steps below.

Step 2902 includes receiving from a first entity, via a receiver, a first medical scan and a medical report corresponding to the medical scan. Step 2904 includes identifying a set of patient identifiers in a subset of fields of a first header of the first medical scan. Step 2906 includes performing a header anonymization function on each of the set of patient identifiers to generate a corresponding set of anonymized fields. Step 2908 includes generating a first de-identified medical scan by replacing the subset of fields of the first header of the first medical scan with the corresponding set of anonymized fields. Step 2910 includes identifying a first subset of patient identifiers of the set of patient identifiers in the medical report by searching text of the medical report for the set of patient identifiers. Step 2912 includes performing a text anonymization function on the first subset of patient identifiers to generate corresponding anonymized placeholder text for each of the first subset of patient identifiers. Step 2914 includes generating a de-identified medical report by replacing each of the first subset of patient identifiers with the corresponding anonymized placeholder text. Step 2916 includes transmitting, via a transmitter, the de-identified first medical scan and the de-identified medical report to a second entity via a network.

In various embodiments, the medical scan is received from a Picture Archive and Communication System (PACS), where the medical report is received from a Radiology Information System (RIS), and where the first de-identified medical scan and the de-identified medical report are transmitted to a central server that is not affiliated with the PACS or the RIS. In various embodiments, first medical scan and the medical report are stored in a first memory for processing. The first memory is decoupled from the network to prevent the set of patient identifiers from being communicated via the network. The first de-identified medical scan and the de-identified medical report are stored in a second memory that is separate from the first memory. The first de-identified medical scan and the de-identified medical report are fetched from the second memory for transmission to the second entity.

In various embodiments, the header anonymization function performed on each of the set of patient identifiers is selected from a plurality of header anonymization functions based on one of a plurality of identifier types of the corresponding one of the subset of fields. In various embodiments, the plurality of identifier types includes a date type. A shift function corresponding to the date type is performed on a first date of the first header to generate the first de-identified medical scan, where the shift function includes offsetting the first date by a determined amount. A second medical scan is received, via the receiver, that includes a second header. A unique patient ID of the first header matches a unique patient ID of the second header. The shift function is performed on a second date of the second header by offsetting the second date by the determined amount to generate a second de-identified medical scan. The second de-identified medical scan is transmitted to the second entity via the network.

In various embodiments, the plurality of identifier types includes a unique patient ID type. A hash function corresponding to the unique patient ID type is performed on the unique patient ID of the first header to generate the first de-identified medical scan. The hash function is performed on the unique patient ID of the second header to generate the second de-identified medical scan. An anonymized unique

patient ID field of the first de-identified medical scan matches an anonymized unique patient ID field of the second de-identified medical scan as a result of the unique patient ID of the first header matching the unique patient ID of the second header.

In various embodiments, the plurality of identifier types includes a linking identifier type that maps the medical scan to the medical report. A hash function corresponding to the linking identifier type is performed on a linking identifier of the first header to generate a hashed linking identifier. A linking identifier field of the first de-identified medical scan includes the hashed linking identifier. Performing the text anonymization function on the first subset of patient identifiers includes determining one of the first subset of patient identifiers corresponds to linking identifier text and performing the hash function on the one of the first subset of patient identifiers to generate the hashed linking identifier, where the de-identified medical report includes the hashed linking identifier.

In various embodiments, a second subset of patient identifiers of the set of patient identifiers is identified in a set of regions of image data of the medical scan by performing an image analysis function on image data of the medical scan. The image analysis function includes searching the image data for the set of patient identifiers. An identifier type is determined for each of the second subset of patient identifiers. One of a plurality of image fiducials is selected for each of the second subset of patient identifiers based on the identifier type. De-identified image data is generated, where a set of regions of the de-identified image data, corresponding to the set of regions of the image data, includes the one of the plurality of image fiducials to obfuscate each of the second subset of patient identifiers. Generating the first de-identified medical scan further includes replacing the image data of the medical scan with the de-identified image data.

In various embodiments, a new patient identifier is identified in the medical report by performing a natural language analysis function on the medical report, where new patient identifier is not included in the set of patient identifiers. The set of patient identifiers is updated to include the new patient identifier prior to searching the image data of the medical scan for the set of patient identifiers, and the second subset of patient identifiers includes the new patient identifier.

In various embodiments, the memory further stores a global identifier blacklist. The natural language analysis function includes searching the medical report for a plurality of terms included in the global identifier blacklist to identify the new patient identifier. In various embodiments, the de-identification system determines that the global identifier blacklist does not include one of the set of patient identifiers, and the global identifier blacklist is updated to include the one of the set of patient identifiers.

In various embodiments, performing the image analysis function further includes identifying a new patient identifier in the image data, where new patient identifier is not included in the set of patient identifiers. Identifying text is extracted from a region of the image data corresponding to the new patient identifier. The new patient identifier is identified in the medical report by searching text of the medical report for the identifying text. The text anonymization function is performed on new patient identifier to generate anonymized placeholder text for the new patient identifier. Generating the de-identified medical report further includes replacing the identifying text with the anonymized placeholder text for the new patient identifier.

In various embodiments, generating the de-identified image data further includes detecting an identifying facial structure in the image data of the medical scan. Generating the de-identified image data includes performing a facial structure obfuscation function on the image data, and where the de-identified image data does not include the identifying facial structure.

FIGS. 12A-12G illustrate an embodiment of a medical scan hierarchical labeling system 3002. The medical scan hierarchical labeling system 3002 can be utilized to generate structured labeling data for medical scans via one or more client devices 120, based on user input to an interactive interface displayed on a display device corresponding to the one or more client devices.

As shown in FIGS. 12A-12G, the medical scan hierarchical labeling system 3002 can communicate bi-directionally, via network 150, with the medical scan database 342, medical scan analysis function database 346, and/or other databases of the database storage system 140, with one or more client devices 120, and/or, while not shown in FIG. 12A, one or more subsystems 101 of FIG. 1. In some embodiments, the medical scan hierarchical labeling system 3002 is an additional subsystem 101 of the medical scan processing system 100, implemented by utilizing the subsystem memory device 245, subsystem processing device 235, and/or subsystem network interface 265 of FIG. 2A. In some embodiments, the medical scan hierarchical labeling system 3002 is implemented by utilizing, or otherwise communicates with, the central server 2640. For example, some or all of the databases of the database storage system 140 are populated with de-identified data generated by the medical picture archive integration system 2600. In some embodiments, the medical scan hierarchical labeling system 3002 can receive de-identified medical scans, annotation data, and/or reports directly from the medical picture archive integration system 2600. For example, the medical scan hierarchical labeling system 3002 can request de-identified medical scans, annotation data, and/or reports that match requested criteria, for example, corresponding to training set criteria. In some embodiments, some or all of the medical scan hierarchical labeling system 3002 is implemented by utilizing other subsystems 101 and/or is operable to perform functions or other operations described in conjunction with one or more other subsystems 101. In some embodiments, the medical scan hierarchical labeling system 3002 is integrated within and/or utilizes the medical scan assisted review system 102 and/or the medical scan annotator system 106.

As shown in FIG. 12A, the medical scan hierarchical labeling system 3002 can store labeling application data 3020 of a labeling application associated with the medical scan hierarchical labeling system 3002. The labeling application data 3020 can include a plurality of prompt decision trees. The plurality of prompt decision trees can include a diagnosis prompt decision tree 3022, a characterization prompt decision tree 3024, and/or a localization prompt decision tree 3026. The labeling application data 3020 can be sent to the one or more client devices. The labeling application data can be stored as a client application as illustrated in FIG. 2A, stored in client memory device 240. The client processing device 230 can execute operational instructions of the labeling application data to enable the corresponding client device 120 to run the labeling application. The labeling application can include an interactive interface 3075 displayed on display device 270.

The client device can further receive a medical scan from the medical scan for labeling, for example, as a transmission

from the medical scan hierarchical labeling system 3002, fetched directly from the medical scan database 342, and/or uploaded to the client device directly. As shown in FIG. 12B, the client device can utilize the labeling application to generate labeling data for the medical scan for transmission back to the medical scan hierarchical labeling system 3002. The labeling data can be mapped to the medical scan in the medical scan database, and can correspond to some or all fields of a medical scan entry 352, such as diagnosis data 440.

The medical scan database can correspond to a relational database and/or a database with a structured set of fields for the plurality of medical scan entries. The labeling data generated via the labeling application can correspond to the structured set of fields. In particular, each of the plurality of prompt decision trees can include leaf nodes that correspond to the structured set of fields, and the labeling application can display prompts to a user of the client device in accordance with the plurality of prompt decision trees. The labeling application can generate labeling data that corresponds to leaf nodes of the plurality of prompt decision tree, based on user input to the prompts indicating selections from sets of selection options that correspond to subsets of the structured set of fields. Requiring that all labeling data adheres to a uniform structure with a discrete set of possibilities in this fashion allows the labeling data to be consumable and easily utilized by other systems. For example, the labeling data can be utilized as training data for one or more other subsystems 101 train a medical scan analysis function.

The labeling application can be utilized by users such as radiologists and/or other labelers responsible for labeling and/or annotating medical scans of the medical scan database with diagnosis data. The interactive interface 3075 can display image data of the medical scan to such a user in conjunction with a plurality of prompts to provide diagnosis, characterization, and/or localization labeling of the medical scan for at least one abnormality identified by the user. The plurality of prompts can present a fixed set of differential diagnosis options allowing the user to select one or more of the differential diagnosis options corresponding to one or more abnormalities detected by the user in the medical scan. The plurality of prompts can present a fixed set of characterization options to classify, describe, or otherwise characterize the one or more differential diagnoses identified in the fixed set of diagnosis data options. The plurality of prompts can present a fixed set of localization options to indicate a region of interest and/or specific anatomical location associated with the one or more differential diagnoses identified in the fixed set of diagnosis data options.

The fixed sets of diagnosis options, characterization option, and/or localization options can correspond to a fixed set of hierarchical options, and can be characterized by a diagnosis prompt decision tree, a characterization prompt decision tree, and/or a localization prompt decision tree, respectively. FIGS. 12C, 12D, and 12E illustrate examples of a diagnosis prompt decision tree 3022, a characterization prompt decision tree 3024, and/or a localization prompt decision tree 3026, respectively. As illustrated, the diagnosis prompt decision tree 3022, the characterization prompt decision tree 3024, and/or the localization prompt decision tree 3026 can each include a root node, a plurality of internal nodes that each branch from the root node or another internal node, and a plurality of leaf nodes that each branch from the root node or an internal node. In particular, all of the fixed set of hierarchical diagnosis options, characterization options, and/or localization options can be represented by all

of the leaf nodes of the diagnosis prompt decision tree, the characterization prompt decision tree, and/or the localization prompt decision tree, respectively. Each root node and internal node can include any number of branches corresponding to a plurality of options that can be selected in response to a prompt corresponding to the node. Each node can be located in one of a plurality of levels of the corresponding prompt decision tree, where each level is characterized by a number of branches from the root node.

The plurality of prompt decision trees can be stored in the application data as a data structures and/or abstract data type corresponding to trees in accordance with the structure of the corresponding decision trees. In some embodiments, other data structures and/or abstract data types are employed that indicate the plurality of decision trees and/or that cause the labeling application to select a set of prompts that are presented to the user as dictated by a corresponding prompt decision tree, to select a corresponding set of options for each prompt as dictated by a corresponding prompt decision tree, and/or to select an ordering of the set of prompts as dictated by a corresponding prompt decision tree, based on each user selection from the selected set of options. This can include filtering a total plurality of prompts and/or a total plurality of options based on user selections, in accordance with the corresponding decision tree. Thus, the labeling application can utilize a prompt selection algorithm that presents the prompts and options in accordance with at least one corresponding decision tree, even if the application data does not include any tree data structures. The examples presented in FIGS. 12C, 12D, and 12E present the behavior exhibited in execution of the labeling application by illustrating the prompt selection algorithm for selecting for the set of prompts that are presented to the user, the corresponding set of options for each prompt, and the ordering of the set of prompts. While the organization of data that includes these prompts and/or options can be stored in a corresponding tree structure in the application labeling data, any other data structure can be employed to organize storage of these prompts and/or options. Furthermore, any other data structure can be employed to organize storage of the instructions for such a prompt selection algorithm for selecting the prompts, the ordering of the prompts, and/or options associated with the prompts as discussed herein.

In presenting these prompt selection algorithms, FIGS. 12C-12E illustrate the distinction between different internal nodes and leaf nodes in accordance with the unique path from the root prompt. For example, diagnosis prompt 1.2 corresponds to the internal node reached after selecting selection option 1 from a set of selection options 1-N from the diagnosis root prompt to reach characterization prompt 1, and then selecting selection option 1.2 from a set of selection options 1.1-1.M. As presented in FIGS. 12C-12E, M, N, and R correspond to any integer number of branches from the corresponding node. The use of ellipses “[. . .]” as presented in FIGS. 12C-12E indicates that any number of additional selections were made to reach the corresponding node, and that the corresponding node can be at any level of the prompt decision tree. For example, leaf node 1.M.[. . .] indicates a leaf node that extends any number of branches from prompt 1.M. The prompt decision trees described herein can be of the same or different configurations presented in FIGS. 12C-12E.

The labeling application can require that a leaf node is reached in each prompt decision tree presented to the user, for example, where the interactive interface will not advance to a next prompt decision tree and/or will not exit until the user continues to make selections to ultimately advance to a

leaf node. This ensures that the user's annotation of the medical fully characterizes at least one abnormality of the medical scan, while still producing structured, consumable labeling data.

5 Some or all of nodes of one or more prompt decision trees can correspond to a prompt presented to a user via the interactive interface, where a selected branch from the node to a next node away from the root node is determined based on user input corresponding to selection of an option corresponding to the branch. Furthermore, some or all nodes of one or more prompt decision trees can correspond to options presented to the labeling application internally, where selected branches from these nodes can be automatically determined for example, based on classifier data 420 of the scan itself such as the modality of the scan and/or anatomical region of the scan; based on leaf nodes reached in of other prompt decision trees; and/or based on another automatic determination that does not correspond to user input to the interactive interface. This can be utilized to advance from the root node to an internal node automatically, where a user is presented with the plurality of prompts of a prompt decision tree starting from this automatically selected internal node. This can also be utilized to automatically advance to an internal node or leaf node based on automatic determinations corresponding to selections of branches to advance down the decision tree, where the automatic determinations are made by the labeling application without user input.

The labeling application can include a plurality of diagnosis prompt decision trees, a plurality of characterization prompt decision trees, and a plurality of localization prompt decision trees for each modality, each anatomical regions, and/or for other scan classifier data 420. For example head CTs, chest x-rays, and chest CTs can each correspond to their own diagnosis prompt decision tree, a characterization prompt decision tree, and localization prompt decision tree. In such embodiments, the labeling application can automatically determine a modality, anatomical region, and/or other category of the medical scan, and can further automatically determine which of the plurality of trees will be used based on the determined category.

The interactive interface can display each of a plurality of prompts to a user of the client device, one at a time, in accordance with the diagnosis prompt decision tree, characterization prompt decision tree, and/or localization prompt decision tree of the labeling application data 3020. The user can select one of a fixed set of options indicated by each of the plurality of prompts of the diagnosis prompt decision tree as user input via interaction with the interactive interface 3075. The selected one of the plurality of options can dictate the next one of the plurality of prompts, in accordance with the corresponding prompt decision tree. These prompts can continue to be displayed in sequence to the user, progressing to the next prompt indicated by the prompt decision tree as the user selects one of the plurality of options presented in accordance with each prompt, until a leaf node of the prompt decision tree is reached. The leaf node can indicate some or all of the selections made by the user in previous nodes of the prompt decision tree from a root node to the leaf node.

Furthermore, each set of options presented for a prompt corresponding to an internal node can include only an appropriate set of options that are possible given the previous selections corresponding to branches from the root node to the diagnosis prompt node. This effectively filters the set of options presented to the user given the selections made so far, ensuring the set of presented options includes only valid

options. These appropriate sets of options can be predetermined by an administrator or other user responsible for creating the prompt decision trees. In some embodiments, the medical scan hierarchical labeling system **3002** automatically determines the options presented for each prompt of a decision tree based on a known set of rules, for example, corresponding to possible options given modality, anatomical regions, and user selections that must be made to reach the corresponding node. In some embodiments, the medical scan hierarchical labeling system **3002** automatically optimizes an ordering of some or all of the prompts. In some embodiments, identical prompts are included in multiple paths of a prompt decision tree, but are presented in different orders. In some embodiments, an ordering of prompts to each leaf node is determined to optimize an expected number of branches necessary to reach a leaf node, optimize the average number of branches to reach a leaf node, to reduce or otherwise optimize the number of selection options presented in one or more prompts of one or more nodes, and/or to otherwise optimize the prompt decision tree. In some embodiments, the set of prompts to reach each leaf node is automatically determined to ensure the corresponding abnormality of each leaf node will be fully described in reaching the leaf node. In some embodiments, the number of this set of prompts is minimized, while still ensuring the corresponding abnormality of each leaf node will be described to a necessary extent in reaching the leaf node.

The interactive interface can first present the user with a prompt asking whether or not an abnormality is present, as shown in FIG. **12C**. If the user selects “no”, the diagnosis prompt decision tree can immediately branch to a leaf node indicating no abnormality is present, and the labeling data can correspond to a structured entry of the medical scan database indicating no abnormality is present. If the user selects the option indicating an abnormality is present, the interactive interface can display a diagnosis root prompt with a plurality of selection options 1-N. The user can select a single one of the N selection options to progress to the corresponding next diagnosis prompt 1-N. For example, the user can select option 1, and the interactive interface will display diagnosis prompt 1 in response, indicating a set of selection options 1.1-1.M. If the user selects option 1.1.1 in response to diagnosis prompt 1.1, leaf node 1.1.1 is reached. Labeling data corresponding to the diagnosis can indicate the diagnosis corresponding to leaf node 1.1.1. The leaf node can effectively indicate all of the selections made from the root node to reach the leaf node. Thus, the labeling data generated in response to selection of the leaf node can further indicate each of the individual selections made by the user to reach the leaf node, for example, where the labeling data indicates selection option 1, selection option 1.1, and selection option 1.1.1.

Each selection option 1, 1.1, and 1.1.1 can correspond to its own field of the relational database, and/or the relational database can include fields corresponding to only the set of leaf nodes, for example, with a binary indication of whether or not the diagnosis corresponding to each leaf node is present. In some embodiments, the database includes a single diagnosis field populated by one of the fixed set of options indicated by the one of the plurality of leaf nodes of the diagnosis prompt tree. The medical scan database can employ any structure with a discrete set of possible diagnosis entries corresponding to a corresponding discrete set of diagnosis options corresponding to the discrete set of leaf nodes of the diagnosis decision tree.

In some embodiments, the labeling application can automatically determine a starting node of the diagnosis decision tree, corresponding to one of the interior nodes rather than the root node. In particular, the labeling application can automatically determine the starting node of the diagnosis decision tree based a determined modality of the medical scan, based on a determined anatomical region of the medical scan, and/or other scan classifier data **420** of the medical scan. Leaf nodes branching from the selected starting node will only include diagnosis options corresponding to abnormalities that can be detected in the corresponding anatomical region of the medical scan and/or that can be detected in the corresponding modality of the medical scan.

For example, the labeling application automatically proceed to selection option 1 that indicates plurality of abnormalities associated with the head in response to determining that the medical scan is a head CT scan. In such embodiments, diagnosis prompt 1 will only present options corresponding to abnormalities that correspond to the head, and that can be detected in a head CT. Any leaf nodes 1.[. . .] that branch from diagnosis prompt 1 will only correspond to final diagnosis that correspond to the head, and that can be detected in a head CT. For example, the fixed set of options indicated by leaf nodes 1.[. . .] can include “brain tumor”, and “brain bleed”, but will not include “wrist fracture” or “pulmonary embolism.”

As another example, the labeling application can proceed automatically to a selection option 2.1 in response to determining that the medical scan is a chest x-ray. For example, a selection option from the diagnosis root prompt can correspond to selecting the anatomical region of the medical scan from a set of anatomical regions, where selection option 2 corresponds to the chest. Diagnosis prompt 2 can correspond to selecting the modality from a plurality of modalities, where selection option 2.1 corresponds to an x-ray. In such embodiments, diagnosis prompt 2.1 will only present options corresponding to abnormalities that correspond to the chest, and that can be detected in an x-ray of the chest. Any leaf nodes 2.1.[. . .] that branch from diagnosis prompt 2.1 will only correspond to final diagnosis that correspond to the chest, and can be detected in a chest x-ray. For example, the fixed set of options indicated by leaf nodes 2.1.[. . .] can include “rib fracture” or “pneumonia” because they can be identified by a chest x-ray. However, “pulmonary embolism” is not included in the leaf nodes 2.1.[. . .], even though this condition is associated with the chest, because leaf nodes 2.1.[. . .] correspond to an x-ray modality, and because reviewing an x-ray alone is not sufficient to determine whether a pulmonary embolism is present or absent. However, if selection option 2.2 corresponds to a CT scan, any leaf nodes 2.2.[. . .] that branch from diagnosis prompt 2.2 will only correspond to a final diagnosis that corresponds to the chest and that can be detected in a chest CT. Thus, leaf nodes 2.2.[. . .] can include “pulmonary embolism” as a selection option because this condition can be determined to be present or absent in reviewing a chest CT.

In some embodiments, a medical scan entry can include any variable number of differential diagnosis, where each of the variable number of differential diagnosis are included in the same field or the variable number of different fields, and where each of the variable number of differential diagnosis corresponds to a different one of the discrete set of leaf nodes of the diagnosis decision tree. In such embodiments, after reaching a leaf node of the diagnosis decision tree, the user can be prompted with a question of whether or not there is another abnormality to characterize. In response to the

user selecting “yes,” the interactive interface can return to the root prompt, and/or the starting node determined based on the anatomical region and/or modality. Furthermore, a deeper starting node, branching from the initial starting node but further away from the root node, can be selected automatically as a result of selection of the first one or more leaf nodes that have already been selected, if these first one or more leaf nodes can narrow the set of options possible for additional differential diagnoses. In some embodiments, characterization and/or localization of the abnormality indicated in the first leaf node is completed in accordance with the characterization prompt decision tree and/or the localization prompt decision tree first, and once this characterization and/or localization is completed, the user is presented the option to identify, characterize, and/or localize additional abnormalities as differential diagnoses.

Once a leaf node is reached in the diagnosis decision tree and/or once it is indicated that no further abnormalities are present, the interactive interface can advance to prompts of the characterization prompt decision tree. FIG. 12D illustrates an example embodiment of a characterization prompt decision tree. The interactive interface can begin with the root node of the characterization prompt decision tree and/or the labeling application can automatically advance to an internal node based on the anatomical region and/or modality as discussed in conjunction with the diagnosis prompt decision tree. Furthermore, as the set of options that can be used to characterize an abnormality depends on the type of abnormality itself, the labeling application can automatically advance to an even deeper internal node based on the leaf node of the diagnosis decision tree that was ultimately selected. For example, as shown in FIG. 12C, a characterization prompt X progressing from leaf node 1.1.1 can correspond to any characterization prompt node of the characterization prompt decision tree determined to be an automatic starting node given a diagnosis corresponding to leaf node 1.1.1. Other leaf nodes of the diagnosis prompt decision tree can also cause the interactive interface to automatically progress to their own pre-determined starting node of the characterization prompt decision tree.

From the selected starting node of the characterization prompt decision tree, the corresponding diagnosis can be characterized based on selections from the starting node, where the interactive interface presents successive prompts corresponding to the nodes branching based on selections to previous prompts, until a leaf node of the characterization prompt decision tree is reached. In embodiments where multiple leaf nodes for multiple abnormalities have already been selected from the diagnosis prompt decision tree, each of these abnormalities can be characterized separately, each with their own automatically determined starting node based on the corresponding leaf node, where a separate leaf node of the characterization decision tree is ultimately reached for each of these differential diagnoses.

The user can indicate if one or more differential diagnoses, identified via selection of multiple leaf nodes of the diagnosis prompt decision tree, correspond to side effects of a main one of the differential diagnoses. This can be performed in the plurality of characterization prompts, where one or more side effects are indicated and characterized as part of characterizing the main diagnosis. As one example, illustrated in FIG. 12D, once a diagnosis is characterized and a leaf node is reached, the user will be prompted with a question asking whether or not the diagnosis is further characterized by any side effects. If so, a selected starting diagnosis prompt Y of the diagnosis prompt decision tree will be presented, and the interactive interface will present

prompts of the diagnosis prompt decision tree until a leaf node of the diagnosis prompt decision tree is ultimately selected for each side effect. Furthermore, each side effect can be further described via selection of leaf nodes in the characterization prompt decision tree and/or the localization prompt decision tree for the each side effect. As another example, after the user identifies a plurality of differential diagnoses, the interactive interface can present a prompt to select at least one main diagnosis from the plurality of differential diagnoses, and can further present a prompt to identify at least one of the remaining differential diagnoses as side effects of the main diagnosis.

For example, if a head CT includes a brain tumor, brain bleed, and fracture, the brain bleed and fracture can be indicated as side effects of the brain tumor. The user can reach leaf nodes of the diagnosis decision tree corresponding to each of “brain tumor”, “brain bleed”, and “fracture”. In particular, after the “brain tumor” leaf node is reached in the diagnosis prompt decision tree, the user can be prompted to identify one or more side effects of the brain tumor as part of characterizing the brain tumor via the characterization prompt decision tree. In response to the user selecting to identify a side effect, prompts of the diagnosis prompt decision tree will be displayed until the leaf node corresponding to “brain bleed” is ultimately reached. In response to the user selecting to identify an additional side effect, prompts of the diagnosis prompt decision tree will be displayed until the leaf node corresponding to “fracture” is ultimately reached. In response to the user selecting that no more side effects are present, the user can continue characterizing and/or advance to localizing the brain tumor itself. The brain bleed and fracture can also each be characterized and localized by ultimately reaching leaf nodes of the characterization prompt decision tree and localization prompt decision tree for each of these side effects.

Once a leaf node is reached in the characterization decision tree and/or once the user indicates that no further side effects are present, the interactive interface can advance to prompts of the localization prompt decision tree 3026, presented in FIG. 12E. The interactive interface can begin with the root node of the localization prompt decision tree and/or the labeling application can automatically advance to an internal node based on the anatomical region and/or modality as discussed in conjunction with the diagnosis prompt decision tree and the characterization prompt decision tree. Furthermore, as the set of options that can be used to localize an abnormality can depend on the type of abnormality itself, the labeling application can automatically advance to an even deeper internal node based on the leaf node of the diagnosis prompt decision tree that was ultimately selected, and/or based on the leaf node of the characterization prompt decision tree that was ultimately selected. For example, as shown in FIG. 12D, a localization prompt Z progressing from leaf node 1.2.R can correspond to any localization prompt node of the localization prompt decision tree determined to be an automatic starting node given a diagnosis and/or characterization corresponding to leaf node 1.2.R. Other leaf nodes of the characterization prompt decision tree can also cause the interactive interface to automatically progress to their own pre-determined starting node of the localization prompt decision tree.

From the selected starting node of the localization prompt decision tree, the corresponding abnormality can be localized based on selections from the starting node, where the interactive interface presents successive prompts corresponding to the nodes branching based on selections to previous prompts, until a leaf node of the localization

prompt decision tree is reached. In embodiments where multiple leaf nodes for multiple abnormalities have already been selected from the diagnosis prompt decision tree as differential diagnoses and/or side effects, each of these abnormalities can be localized separately, each with their own automatically determined starting node based on the corresponding leaf node, where a separate leaf node of the localization decision tree is ultimately reached for each of these abnormalities.

In some embodiments, as an abnormality might be located in multiple places, the user can select any number of leaf nodes to localize each abnormality. For example, as shown in FIG. 12E, after reaching a leaf node such as leaf node 1.1.1, the user can be presented with an option to identify other locations in which the abnormality is present. In response to the user selecting "yes," the interactive interface can return to the root prompt, and/or the starting node determined based on the anatomical region and/or modality. Furthermore, a deeper starting node, branching from the initial starting node but further away from the root node, can be selected automatically as a result of selection of the first one or more leaf nodes of the localization prompt decision tree that have already been selected, if these first one or more leaf nodes can narrow the set of options possible for additional locations diagnoses, such as narrowing to a set of adjacent locations or locations within proximity of locations selected thus far.

At least one internal node of the localization prompt decision tree can require that multiple leaf nodes are reached to characterize the abnormality. For example, as shown in FIG. 12G, localization prompt 1 indicates that a number Q leaf nodes are required. Once a leaf node, such as leaf node 1.2.1 is reached, if the required number Q leaf nodes have not yet been selected, localization prompt 1 will be presented and the user can progress to leaf nodes, returning to localization prompt 1 until the required number of leaf node selections Q have been made. While the diagnosis prompt decision tree of FIG. 12C nor the characterization prompt decision tree of FIG. 12D illustrate internal nodes or root nodes requiring multiple leaf nodes be reached, such nodes requiring multiple leaf node selections can similarly be included in the diagnosis prompt decision tree or the characterization prompt decision tree in other embodiments.

In some embodiments, this required number of leaf nodes can inherently be built into the localization prompt decision tree with single required selection from each node. For example, instead of returning to localization prompt 1 from leaf node 1.2.1, the leaf node 1.2.1 can instead be an internal node 1.2.1 presenting localization prompt 1, with the same set of M selection options this internal node, allowing the user to continue to eventually reach the Q leaf node selections as a result of progressing down deeper sets of branches and identifying each of the Q leaf nodes as selections to the final, Qth leaf node along the way.

As a particular example, if a leaf node corresponding to "brain tumor" is selected from the diagnosis prompt decision tree, at least one lobe and at least one compartment must be selected in the step of localizing the brain tumor. The localization prompt decision tree can present the user with a plurality of lobe selection options and a plurality of compartment selection options, and can require that at least one leaf node corresponding to a lobe is selected, and that at least one leaf node corresponding to a compartment is selected. For example, once a lobe is selected, the same prompt can be presented, excluding the lobe that has already been selected, until one of the plurality of compartment options is also selected. As another example, the localization prompt

decision tree can present a first prompt with only a plurality of lobe options in an internal node, and once one of the plurality of lobe options is selected, can branch to the next node that presents a second prompt with only a plurality of compartment options, where the leaf node ultimately reached indicates the selected lobe and the selected compartment. In such embodiments, each of the plurality of lobe selection options can branch to the identical prompt nodes presenting the plurality of compartment options.

In some embodiments, based on the diagnosis leaf node, characterization leaf node, and/or localization leaf node, an additional set of questions may be necessary. For example, when leaf node N. [. . .] is reached, at least one additional prompt will be presented. This additional set of questions can correspond to prompts of at least one additional prompt decision tree, where the questions are hierarchical and dependent on previous questions, and/or where leaf nodes of each additional prompt decision tree must be reached. The additional prompt set can be the same or different for different leaf nodes. The additional prompt set can depend only on the diagnosis leaf node, characterization leaf node, localization leaf node, and/or on a combination of two or more nodes. The additional prompt set can be further based on differential diagnoses or side effects. The additional prompt set can be based on the modality of the medical scan and/or the anatomical region of the medical scan. For example, if the medical scan corresponds to a head CT, the additional question set includes ventricular system questions that will automatically be presented via the interactive interface in response to an automatic determination that the medical scan is a head CT. In some embodiments, the additional prompt set is inherently included as additional internal nodes of the diagnosis prompt decision tree, characterization prompt decision tree, and/or localization prompt decision tree, ultimately reaching a final leaf node that includes answers to the additional questions. The answers to the additional questions can be included in the fixed format of the labeling data, where each of the additional questions similarly presents a fixed set of options. Alternatively or in addition, at least one additional question can correspond to unstructured data, such as text or voice input by the user, drawings and/or or shapes outlining one or more abnormalities superimposed upon the medical scan as user input entered by the user, measurement data indicating size, shape, diameter, and/or volume of abnormalities as identified by the user, a report entered by the user, or other unstructured data. In some embodiments, this unstructured data can be mapped to the medical scan in the database, but can be separate from the rest structured labeling data.

In some embodiments, the medical scan includes a plurality of image slices. The user can be prompted to select a proper subset of slices and/or a single slice that includes an abnormality to be described in the labeling data. A user selection that indicates a selected subset of the plurality of images slices of the medical scan can be received via user input to the interactive interface. This prompt can be included in the localization prompt decision tree and can dictate further prompts of the localization prompt decision tree based on a narrowed anatomical location corresponding to the proper subset of slices and/or the single slice. Alternatively, this prompt can be presented separately from prompts of the localization prompt decision tree, and the starting localization prompt can be selected based on the selected subset of the plurality of image slices. For example, if a subset of slices selected by the user indicates the frontal

lobe of a head CT, localization prompts presented by the user interface will include options corresponding to the selection of the frontal lobe.

The user can be prompted to provide an urgency ranking as part of a prompt decision tree, additional question set, and/or as a final prompt presented to the user. The urgency ranking prompt can similarly include a fixed set of urgency ranking options for selection by the user to indicate an urgency associated with the diagnosis, associated with further review of the medical scan, and/or associated with further scans, tests, or appointments with the patients that may be necessary. The urgency ranking can be included in the labeling data, and/or can be utilized in triaging of the medical scan by one or more subsystems.

In some embodiments, one or more of the nodes of one or more of the prompt decision trees can be optional, where a user selection is not required. Such nodes can include a "skip" branch, indicating that selection of one or more of the selection options is not necessary for this prompt. Selecting the skip option can correspond to a branch that advances to a next node in the prompt decision tree.

While the discussion thus far indicates that the user first selects a leaf node of the diagnosis prompt decision tree, then a leaf node of the characterization prompt decision tree, and finally a leaf node of the localization prompt decision tree, prompts of the diagnosis prompt decision tree, the characterization prompt decision tree, and the localization prompt decision tree can be presented in any order. For example, the user might first localize the abnormality, and then based on this localization, a starting node of the diagnosis prompt decision tree is determined based on types of abnormalities that can exist and/or be observed at the particular location indicated in the localization leaf node.

Furthermore, while FIGS. 12C-12E present the diagnosis prompt decision tree, characterization prompt decision tree, and localization prompt decision tree separately, prompts corresponding to diagnosis, characterization, and localization of one or more abnormalities can be included in a single prompt decision tree of the labeling application data, where reaching a leaf node of the single prompt decision tree includes selecting at least one diagnosis option from at least one internal node corresponding to a diagnosis prompt, selecting at least one characterization option from at least one internal node corresponding to a characterization prompt, and selecting at least one localization option from at least one internal node corresponding to a localization prompt. In some embodiments, the flow of prompts to at least one leaf node includes a plurality of diagnosis prompts, characterization prompts, and localization prompts in any order, for example, where first a localization prompt is presented, then a diagnosis prompt, then another localization prompt, and then a characterization prompt. In some embodiments, multiple prompt decision trees that includes this mix of diagnosis prompts, characterization prompts, and/or localization prompts are included in the labeling application data. Each of these multiple prompt decision trees can correspond to a different modality, a different anatomical region, a different modality/anatomical region pair, and/or other different scan classifier data 420.

The labeling application can be utilized by multiple client devices corresponding to multiple users, and each user can label multiple medical scans by utilizing the labeling application. Labeling data generated over time by one or more users for one or more medical scans can be tracked in a user database 344, where the number of scans labeled by each user of each client device is tracked. The user database can further track how many scans have been labeled for each of

a set of scan categories corresponding to different modalities, different anatomical regions, and/or other categories indicated by scan classifiers data 420. Users can be incentivized and/or rewarded for reaching a threshold number of labeled scans in one or more scan categories and/or reaching a total number of labeled scans. Similarly, users can be incentivized and/or rewarded for reaching and/or maintaining a threshold labeling rate in one or more scan categories and/or reaching and/or maintaining a threshold labeling rate overall. In some embodiments, users can be incentivized to label scans in categories with a count and/or labeling rate that is below a threshold for that user, encouraging the user to expand their skills and label scans of modalities and/or anatomical regions they do not typically label. In some embodiments, users can be incentivized to label scans in categories with a count and/or labeling rate that is below a global threshold across all users, for example, corresponding to scan types that are not addressed enough by users across the system, to encourage the user to meet this need. The incentives can include financial incentives, can include a favorable adjustment to performance score data and/or qualification data, and/or can include advancing a user to an expert status in one or more scan categories in which a number and/or labeling rate compares favorably.

In some embodiments, a medical scan can be automatically pre-processed to partition the medical scan in accordance with multiple anatomical regions included within the medical scan. For example, a full body scan can be partitioned into a set of medical scan portions, where each medical scan portion corresponds to each of a set of anatomical regions. For example, the full body scan can be partitioned into medical scan portions corresponding to the head, chest, arm, leg, etc. for individual labeling. These partitions can each be labeled by the same user or by different users. For example, the medical scan hierarchical labeling system 3002 can perform this pre-processing step prior to transmission of the medical scan to a client device. The different medical scan portions can be sent to different users for labeling based on determining each user has favorable qualification data and/or performance score data for the corresponding anatomical region. The labeling data can be retrieved from all of the users and can be compiled for the original medical scan to be mapped to the medical scan database. In some embodiments, the pre-processing step is performed after a medical scan is retrieved by a client device as part of execution of the labeling application. Each partition can be presented in conjunction with prompt decision trees corresponding to the anatomical region of each partition and/or in conjunction with starting nodes of the prompt decision trees determined based on the anatomical region of each partition.

As presented in FIGS. 12A and 12B, the labeling application is executed by the client device, allowing the client device to generate all of the labeling data locally, and this final labeling data is then transmitted back to the medical scan hierarchical labeling system 3002 via the network. In some embodiments, some or all of the steps performed by the client device in accordance with execution with the labeling application can be instead executed by the medical scan hierarchical labeling system 3002. This can be accomplished via additional transmissions between the medical scan hierarchical labeling system 3002 and the client device.

For example, some or all of the user input, such as user selection of a selection option of a prompt, can be transmitted via the network the medical scan hierarchical labeling system 3002. The medical scan hierarchical labeling system 3002 can utilize the corresponding prompt decision tree,

stored in memory of the medical scan hierarchical labeling system **3002**, to determine the next prompt that will be presented to the user and the corresponding set of options. This next prompt and corresponding set of options can be transmitted to the client device for display via the interactive interface, where one of this set of options is selected by the user via user input, and is transmitted back to the medical scan hierarchical labeling system **3002**. Another next prompt and another next set of options is determined by medical scan hierarchical labeling system **3002** based on the prompt decision tree for transmission back to the client device. This process can continue until a leaf node is ultimately selected.

As another example, the medical scan hierarchical labeling system **3002** can automatically select the starting node for one or more of the prompt decision trees based on the anatomical region, modality, and/or other features of the medical scan, in conjunction with transmission of the medical scan. An indicator of the starting node can be transmitted to the client device, and the client device can present the plurality of prompts beginning with the starting node based on the indicator of the starting node received from the medical scan hierarchical labeling system **3002**.

As another example, the medical scan hierarchical labeling system **3002** can automatically select the starting node for one or more of the prompt decision trees based on the anatomical region, modality, and/or other features of the medical scan that is transmitted to the client device. A subset of the corresponding one or more prompt decision trees is selected by the medical scan hierarchical labeling system **3002**, where the root node of the subset is the selected starting node, and where the subset includes internal nodes and leaf nodes that extend from the starting node. This subset of the one or more prompt decision trees can be transmitted to the client device for use, where the remainder of the one or more prompt decision trees is not transmitted to the client device and/or not stored by the client device. This subset of the one or more prompt decision trees can be utilized by the client device to present the plurality of prompts.

As shown in FIG. **12F**, a medical scan training set can be retrieved from the medical scan database **342**. The medical scan hierarchical labeling system **3002**, or another subsystem **101**, can perform a training step **3090**, for example, by utilizing the medical scan image analysis system and performing training step **1352** of FIG. **7A**. The medical scan training set can include medical scans that were labeled by one or more client devices by utilizing the medical scan hierarchical labeling system **3002** as discussed in conjunction with FIGS. **12A-12E**. In particular, labeling data mapped to the medical scans of the medical scan training set was generated by client devices based leaf nodes reached in the diagnosis prompt decision tree, characterization prompt decision tree, and/or localization prompt decision tree. This labeling data is indicated for each medical scan in the medical scan training set, for example, as a structured entry in the medical scan database. The labeling data for each medical scan in the training set can be utilized as an output feature vector and/or output nodes of a neural network in the training step **3090**, where the output feature vector and/or output nodes are structured in accordance with the fixed structure of the labeling data. The input feature vector and/or input nodes of the neural network can correspond to image data of each medical scan in the training set, patient history data of each medical scan in the training set, and/or other data of the corresponding medical scan entry of each medical scan in the training set. Once the model is trained, model data can be transmitted to the medical scan function analysis

database **346** to be utilized by one or more other subsystems **101**. Alternatively or in addition, the model data can be transmitted to the medical picture archive integration system **2600** to be utilized in performance of inference functions by the medical picture archive integration system **2600**.

The model data can be utilized by the medical scan hierarchical labeling system **3002**, or another subsystem **101**, as shown in FIG. **12G** to generate inference labeling data for one or more new medical scans to be labeled, for example, retrieved from the medical scan database **342**, by performing an inference function **3095** that utilizes the trained model. For example, the medical scan image analysis system **112** and/or the inference step **1354**, detection step **1372**, and/or abnormality classification step **1374** of FIG. **7B** can be utilized to generate inference labeling data for new medical scans. The inference labeling data can correspond to the same fixed structure of the labeling data generated by utilizing the labeling application, dictated by the fixed set of abnormality options, the fixed set of classification options, and the fixed set of localization options. The inference labeling data can indicate probability values for one or more of the fixed set of abnormality options, one or more of the fixed set of classification options, and/or one or more of the fixed set of localization options. These probability values can indicate probabilities that one or more of the fixed set of abnormality options is present in the scan and/or in a region of the scan corresponding to a probability matrix corresponding to the probability value, that one or more of the fixed set of classification options describes the one or more of the fixed set of abnormality options, and/or that the abnormality is located in the one or more of the fixed set of localization options. The inference labeling data can be mapped to the new medical scan in the medical scan database and/or can be transmitted to a client device for display to a user via a user interface.

In various embodiments, medical scan hierarchical labeling system includes a medical scan database that stores a plurality of medical scan entries, at least one processor, and a memory. The memory stores labeling application data that includes application operational instructions and a plurality of prompt decision trees. The plurality of prompt decision trees includes a diagnosis prompt decision tree, a characterization prompt decision tree, and a localization prompt decision tree. Each of the plurality of prompt decision trees includes a root node, a set of internal nodes, and a set of leaf nodes. Each root node and each of the set of internal nodes correspond to one of a plurality of prompts. Each root node and each of the set of internal nodes include a set of branches that each correspond to one of a discrete set of selection options for the one of the plurality of prompts. The memory further stores a medical scan relational database that stores a plurality of medical scan entries. The medical scan relational database includes a discrete set of fields corresponding to the leaf nodes of the plurality of prompt decision trees.

The executable instructions, when executed by the at least one processor, cause the medical scan hierarchical labeling system to transmit, via a network, labeling application data to a client device for storage. The application operational instructions of the labeling application data, when executed by at least one client device processor of the client device, cause the client device to execute a labeling application. The executable instructions further cause the medical scan hierarchical labeling system to transmit via the network, a medical scan to the client device.

Execution of the labeling application by the client device causes the client device to, in response to receiving the medical scan, display, via an interactive interface presented

on a display device associated with the client device for display to a user associated with the client device, image data of the medical scan. The client device automatically determines a starting diagnosis prompt by selecting one of the set of internal nodes of the diagnosis prompt decision tree based on an anatomical region of the medical scan and further based on a modality of the medical scan. The client device displays, via the interactive interface, a plurality of diagnosis prompts of the diagnosis prompt decision tree, in succession, beginning with the starting diagnosis prompt, in accordance with corresponding nodes of the diagnosis prompt decision tree until a first one of the set of leaf nodes of the diagnosis prompt decision tree is ultimately selected. The interactive interface progresses to each next one of the plurality of diagnosis prompts by selecting one of the set of branches of each corresponding node in accordance with each of a plurality of corresponding user diagnosis selections, received via user input. Each of the plurality of corresponding user diagnosis selections corresponds to one of the discrete set of selection options for each one of the plurality of diagnosis prompts displayed via the interactive interface.

The client device automatically determines a starting characterization prompt by selecting one of the set of internal nodes of the characterization prompt decision tree based on the anatomical region of the medical scan, based on the modality of the medical scan, and further based on the first one of the set of leaf nodes of the diagnosis prompt decision tree. The client device displays, via the interactive interface, a plurality of characterization prompts, in succession, beginning with the starting characterization prompt, in accordance with corresponding nodes of the characterization prompt decision tree until a first one of the set of leaf nodes of the characterization prompt decision tree is ultimately selected. The interactive interface progresses to each next one of the plurality of characterization prompts by selecting one of the set of branches of each corresponding node in accordance with each of a plurality of corresponding user characterization selections, received via user input. Each of the plurality of corresponding user characterization selections corresponds to one of the discrete set of selection options for each one of the plurality of characterization prompts displayed via the interactive interface.

The client device automatically determines a starting localization prompt by selecting one of the set of internal nodes of the localization prompt decision tree based on the anatomical region of the medical scan, and further based on the modality of the medical scan. The client device displays via the interactive interface, a plurality of localization prompts, in succession, beginning with the starting localization prompt, in accordance with corresponding nodes of the localization prompt decision tree until a first one of the set of leaf nodes of the localization prompt decision tree is ultimately selected. The interactive interface progresses to each next one of the plurality of localization prompts by selecting one of the set of branches of each corresponding node in accordance with each of a plurality of corresponding user localization selections, received via user input. Each of the plurality of corresponding user localization selections corresponds to one of the discrete set of selection options for each one of the plurality of localization prompts displayed via the interactive interface.

The client device transmits, via the network, labeling data that includes a set of labels indicating the first one of the set of leaf nodes of the diagnosis prompt decision tree, the first one of the set of leaf nodes of the characterization prompt

decision tree, and the first one of the set of leaf nodes of the localization prompt decision tree.

The executable instructions, when executed by the at least one processor of the medical scan hierarchical labeling system, further cause the medical scan hierarchical labeling system to receive, via the network, the set of labels from the client device, and to populate a medical scan entry of the medical scan in the medical scan relational database based on the set of labels.

FIGS. 13A-13E presents an embodiment of a medical scan labeling quality assurance system **3004**. The medical scan labeling quality assurance system **3004** can be utilized to routinely gauge how a set of users responsible for labeling medical scans are performing. This can be accomplished by sending a quality assurance set of medical scans to client devices of the set of users, where each user generates labeling data for each of the medical scans in the set. The quality assurance set of medical scans is also sent to a user identified as an expert for labeling. The set of labeling data generated by the expert user, corresponding to labeling data for each medical scan in the set, can be considered a set of golden labeling data. The golden labeling data in this set can be mapped to the corresponding medical scans in the medical scan database. The set of medical scans and corresponding set of golden labeling data can be utilized as training data to train one or more medical scan analysis functions. The set of golden labeling data can be sent to the set of users for review. The set of golden labeling data can be utilized to generate performance score data for the set of labeling data generated by each user. The performance score data can be mapped to the user in the user database and can be utilized to determine whether to dismiss one or more users and/or can be utilized to advance one or more users to an expert status.

As shown in FIGS. 13A-13E, medical scan labeling quality assurance system **3004** can communicate bi-directionally, via network **150**, with the medical scan database **342**, user database **344**, and/or other databases of the database storage system **140**, with one or more client devices **120**, and/or, while not shown in FIG. 12A, one or more subsystems **101** of FIG. 1. In some embodiments, the medical scan labeling quality assurance system **3004** is an additional subsystem **101** of the medical scan processing system **100**, implemented by utilizing the subsystem memory device **245**, subsystem processing device **235**, and/or subsystem network interface **265** of FIG. 2A. In some embodiments, the medical scan labeling quality assurance system **3004** is implemented by utilizing, or otherwise communicates with, the central server **2640**. For example, some or all of the databases of the database storage system **140** are populated with de-identified data generated by the medical picture archive integration system **2600**. In some embodiments, the medical scan labeling quality assurance system **3004** can receive de-identified medical scans, annotation data, and/or reports directly from the medical picture archive integration system **2600**. For example, the medical scan labeling quality assurance system **3004** can request de-identified medical scans, annotation data, and/or reports that match requested criteria, for example, corresponding to training set criteria. In some embodiments, some or all of the medical scan labeling quality assurance system **3004** is implemented by utilizing other subsystems **101** and/or is operable to perform functions or other operations described in conjunction with one or more other subsystems **101**. In some embodiments, the medical scan labeling quality assurance system **3004** is integrated within and/or utilizes the medical scan assisted review system **102** and/or the medical

scan annotator system **106**. In some embodiments, the medical scan labeling quality assurance system **3004** is integrated within and/or utilizes the medical scan hierarchical labeling system **3002**.

In some embodiments, the medical scan labeling quality assurance system **3004** is utilized in conjunction with one or more other subsystems another subsystem **101** responsible for sending and/or triaging medical scans to users for labeling, and mapping the generated labeling data to the medical scans in the medical scan database. The normal operation of this overarching system can include triaging or otherwise assigning each of a plurality of medical scans to one of a set of users in the system for labeling. Thus, labeling data generated for each these medical scan is generated by only one user that was assigned to the medical scan. This labeling data can be mapped to the medical scans in the medical scan database and/or can be used as training data for to train a computer vision model and/or to train one or more medical scan analysis functions. This labeling data generated in normal operation of the system may not be reviewed by other users in the system, unless the user is later flagged as poor performing in conjunction with the quality assurance process.

On the contrary, the set of medical scans selected in conjunction with the quality assurance process are sent to multiple users that generate labeling data, such as all of the users in the system. Labeling data will be generated for this set of medical scans by all of the users for comparison to the golden labeling data generated for these medical scans by the expert user. This quality assurance set of medical scans can comprise only a small subset of the plurality of medical scans reviewed by the set of users. For example, a user may label 100 medical scans a week, where 10 of these medical scans are sent every week in conjunction with the quality assurance process. Labeling data generated by the for these 10 medical scans of the quality assurance set are compared to the golden labeling data to generate performance score data as described herein. Labeling data for the remaining 90 of the medical scans is mapped to the medical scan in the medical scan database normally, without review. Sets of medical scans sent to the set of users within a timeframe corresponding to the schedule of the quality assurance process each include the same quality assurance set of medical scans, as well as additional medical scans that will be processed normally. Thus, the intersection of these sets of medical scans sent to the set of users within the time frame includes only the quality assurance set of medical scans.

In some embodiments, the set of medical scans of the quality assurance set are indistinguishable by users from the remaining medicals scans they receive in the timeframe. This can ensure that users are labeling scans of the quality assurance set normally, and/or can ensure that the labeling data of the quality assurance set is a reflection of normal labeling behavior of the user. Alternatively, the set of medical scans of the quality assurance set can always be the first medical scans sent to the set of users in the timeframe. In some embodiments, the remaining medical scans for normal processing are only sent to the user if the performance score data for the labeling data of the quality assurance set compares favorably to the threshold. This can ensure that users with poor performance are not allowed to generate their own labeling data for medical scans in the medical scan database.

As shown in FIG. **13A**, the medical scan labeling quality assurance system **3004** can initiate performance of the quality assurance process by generating a request for a set of medical scans for transmission the medical scan database,

and by generating request for a set of client device identifiers for a set of users for transmission the user database. The medical scan labeling quality assurance system **3004** can receive a set of medical scans from the medical scan database in response and can receive a set of client device identifiers from the user database. As illustrated in the example presented in FIGS. **13A-13E**, the set of medical scans of the quality assurance process received from the medical scan database includes M medical scans, and the set of client device identifiers corresponding to a set of N users. The M medical scans corresponds to the quality assurance set of medical scans.

As the medical scan labeling quality assurance system **3004** routinely performs a quality assurance process to assess the performance of users at regular intervals, the request for the set of medical scans and request for set of client device identifiers can be transmitted in response to determining a fixed amount of time has elapsed since a last time that the quality assurance process was performed and/or by determining the current time compares favorably to a scheduled time to conduct the quality assurance process. This step can be performed to initiate each subsequent quality assurance process at the routine time intervals and/or in accordance with the subsequently scheduled quality assurance processes.

The medical scan set request can indicate particular identifiers or criteria that the set of medical scans should meet. In some embodiments, a set of medical scans that meet the criteria are randomly or pseudo-randomly selected. In some embodiments, the set of medical scans do not have any diagnosis data **440** or other labeling data associated with them in the medical scan database and/or are otherwise selected in response to determining they need to be labeled. The criteria used to select the set of medical scans can include a number of medical scans to be selected, at least one desired modality, at least one desired anatomical region, other features of scan classifier data **420**, an urgency level, a recency that the medical scans were added to the system, or other criteria. This criteria can be determined via user input by an administrator and/or can be determined automatically by the medical scan labeling quality assurance system **3004**, for example, based on one or more scan categories determined to best assess the set of labelers. For example, in response to determining previous sets of medical scans did not include a threshold number of medical scans corresponding to a scan category, at least one medical scan, or at least a threshold number of medical scans, of that scan category can be included in the set of medical scans.

In some embodiments, types of scans that users perform poorly on can be selected to assess whether users are improving with time and/or to assess whether users should be dismissed from labeling these types of scans in the future. For example, in response to determining one or more of the set of users have performance score data indicating poor performance for a scan category and/or in response to determining the poorest performing scan category across all of the users in the set of users, at least one medical scan, or at least a threshold number of medical scans, of this that scan category can be included in the set of medical scans.

Alternatively, the criteria can be selected to include types of scans that users perform well on can be selected to ensure users are still performing strongly with time and/or to assess whether users should be assigned an expert status for these types of scans. For example, in response to determining one or more of the set of users have performance score data indicating strong performance for a scan category and/or in response to determining the strongest performing scan cat-

egory across all of the users in the set of users, at least one medical scan, or at least a threshold number of medical scans, of this that scan category can be included in the set of medical scans.

A set of client device identifiers for a selected set of users can also be determined, for example, by fetching the identifiers from the user database. The set of users can correspond to users that regularly generate labeling data for medical scans to populate entries in the medical scan database and/or to generate labeling data utilized as training data. For example, the set of users can correspond to all users, or a subset of users, that utilize the labeling application to generate labeling data that is assigned to medical scans in the medical scan database in accordance with the medical scan hierarchical labeling system. The same or different set of users can be selected for each routine quality assurance process over time, based on the same or different criteria determined over time.

In some embodiments, a subset of the set of users is selected based on their corresponding performance history, for example, indicated by their performance score data **530** in the user database **344**. In some embodiments, the set of users can include only to users with performance score data that falls within a determined performance score range. For example, the high end of the performance score range can correspond to a performance score threshold indicating that users meeting or exceeding this threshold need not be assessed at regular intervals. In some embodiments, this high end of the performance score range corresponds to a performance score threshold indicating that users meeting or exceeding this threshold are assigned an expert status. In some embodiments, determining whether performance score data compares favorably to the high end includes determining that the user has labeled at least a threshold number of medical scans and/or at least a threshold proportion of medical scans with a high end performance score. The low end of the performance score range can correspond to a performance score threshold indicating that users failing to meet this threshold are no longer allowed to label medical scans at all, are on probation for a duration of time, are on probation until they can pass a more rigorous test of their labeling capabilities, have been dismissed from the system entirely, or otherwise are not currently labeling any scans for the given timeframe. In some embodiments, the performance score range only has a high end. In some embodiments, the performance score range only has a low end.

The performance score range can be set for one or more scan categories of medical scans in the quality assurance set, and can be the same or different for each category. The medical scan labeling quality assurance system **3004** can select only users that fall within all of the performance score ranges for their corresponding categorized performance scores. Alternatively the medical scan labeling quality assurance system **3004** can select users that fall within at least one of the performance score ranges for their corresponding categorized performance scores.

Performing the quality assurance process can include selecting different quality assurance sets of medical scans for different selected sets of users based on different criteria. For example, a first set of users can be sent only chest x-rays, and a second set of users can be sent head CTs. This grouping of users can be based on determining the types of scans the users specialize in, the types of scans the users labels most frequently, the types of scans the users perform most poorly on, the types of scans the users label least frequently, or other criteria. In some embodiments, some users are sent multiple of these sets of medical scans, where

one user is included in the first set and the second set and is thus sent the chest x-ray quality assurance set as well as the head CT quality assurance set. Each quality assurance set can have golden labeling data generated by the same or different expert user. For example, users identified as experts in each scan category can be sent the corresponding set of medical scans to generate the corresponding set of golden labeling data.

The quality assurance process can be part of an onboarding process for new users to the system, where these new users are not assigned scans for normal labeling until they pass the onboarding process. In such embodiments, these users can receive only the set of quality assurance scans in each timeframe, until their performance score data compares favorably to an onboarding threshold and/or until threshold number of timeframes for which the user participated in the quality assurance process exceeds a threshold. In some embodiments, poor performing users are automatically pushed back to a beginner state and must complete the onboarding process again before continuing to label scans normally.

Once the medical scan set and the set of client devices have been determined for the quality assurance process for the current timeframe, the set of medical scans 1-M can be sent to client devices for each of the set of users 1-N as illustrated in FIG. **13B**. The client devices 1-N can be implemented utilizing client device **120**. The medical scans 1-M can be displayed via interactive interface **3175**. For example, the user can generate labeling data for each medical scan 1-M, in sequence, one at a time, in accordance with display of each medical scan 1-M in sequence, one at a time, by the interactive interface **3175**. The labeling application of the medical scan hierarchical labeling system **3002** can be utilized to enable the user to interact with the interactive interface **3175** to enable each client device generate labeling data for each medical scan. In some embodiments, the labeling data is other annotation data and/or unstructured data discussed herein, and interfaces of the medical scan assisted review system **102** and/or medical scan annotator system **106** can be utilized to enable each client device to generate this unstructured labeling data based on user input. For example, the labeling data can include text entered by the user, can include region of interest data or location data indicating the location of abnormalities identified by the user, can indicate measurement data indicating size, shape, diameter, and/or volume of abnormalities as identified by the user, and/or can include other diagnosis data **440**. The labeling data 1-M generated by each client device for the set of medical scans is transmitted to the medical scan labeling quality assurance system **3004**.

As illustrated in FIG. **13C**, the medical scan labeling quality assurance system **3004** can request a client device identifier for an expert user, and can receive the expert client device identifier in response. The expert user can be selected randomly from a set of expert users, can correspond to a highest ranked user, and/or can otherwise be selected in response to determining the user has an expert status. For example, the expert status can be indicated in the expert data **543** of the user profile entry.

The expert user can be determined to be an expert for a modality, anatomical region and/or other scan classifier data **420** corresponding to at least one of the medical scans of the medical scan set. This determination can be based on categorized performance score data **534** of the user profile entry, categorized qualification data of the user profile entry, and/or another indication that the user is an expert in a category corresponding to the type of some or all of the

medical scans. In some embodiments, a user identified as an expert in a plurality of categories corresponding to a set of categories characterizing all of the set of medical scans is selected, ensuring the expert user is highly qualified to generate golden labeling data for all of the scan types in the set of medical scans. In some embodiments, all of the medical scans in the set are of the same category, and a user identified as an expert of the corresponding category is selected. For example, when a plurality of quality assurance processes corresponding to different scan categories are performed, an expert user can be identified for each of the plurality of quality assurance processes based on having expertise in the corresponding category.

The client device identifier for the expert user can be utilized to send the set of medical scans of the quality assurance process to expert client device **3120**. Expert client device **3120** can be implemented by utilizing a client device **120** that is associated with the expert user. The medical scans 1-M can be displayed via interactive interface **3177**, which can be the same or different from interactive interface **3175**. For example, the expert user can generate golden labeling data for each medical scan 1-M, in sequence, one at a time, in accordance with display of each medical scan 1-M in sequence, one at a time, by the interactive interface **3177**. The labeling application of the medical scan hierarchical labeling system **3002** can be utilized to enable the expert user to interact with the interactive interface **3177** to enable the expert client device generate golden labeling data for each medical scan. In some embodiments, the golden labeling data is other annotation data and/or unstructured data discussed herein, and interfaces of the medical scan assisted review system **102** and/or medical scan annotator system **106** can be utilized to enable the expert client device to generate this unstructured labeling data based on user input. The golden labeling data 1-M generated by the expert client device for the set of medical scans is transmitted to the medical scan labeling quality assurance system **3004**.

The medical scan labeling quality assurance system **3004** can also send each set of labeling data 1-N to the expert client device for review. The labeling data provided by some or all of the users for some or all medical scans can be displayed to the expert user to assist the expert user in creating the global data set. In some embodiments, the expert user can elect to display labeling data from each of the N users one at a time for a single medical scan in conjunction with the display of the single medical scan. In some embodiments, the expert user can elect to display a subset or all of the labeling data from each of the N users one at a time for a single medical scan in conjunction with the display of the single medical scan. The labeling data from the N users can be displayed as text, can outline, highlight or otherwise indicate abnormalities identified in the medical scan by at least one of the N users, or can otherwise indicate the labeling data. If labeling data for multiple users are displayed simultaneously, labeling data of different users can be distinguished, for example, by displaying labeling data in different colors, displaying user identifiers in conjunction with each of the labeling data, or otherwise distinguishing the labeling data created by different users. In some embodiments, historical performance data and/or qualifications for a user can be displayed in conjunction with the labeling data for that user, allowing the expert user to evaluate contrasting opinions based on how well the corresponding users have performed in the past and/or how qualified the corresponding users are.

In some embodiments, a plurality of experts are identified. The set of medical scans are sent to a plurality of client

devices corresponding to the plurality of experts, and the golden labeling data is generated by performing a consensus function on the labeling data received from each of the plurality of experts, such as the annotation consensus function described in conjunction with FIG. **4040**.

The medical scan labeling quality assurance system **3004** can pre-process the labeling data sets 1-N before transmission to the expert client device. The medical scan labeling quality assurance system **3004** can identify the most popular labeling data and/or average labeling data for each medical scan 1-M by evaluating the labeling data from each of the N users for each medical scan. For example, the medical scan labeling quality assurance system **3004** can generate consensus data as discussed in conjunction with the discussion of the medical scan annotator system **106** in FIGS. **14A-14V**. The medical scan labeling quality assurance system **3004** can also identify labeling data from one or more user that differs most from the popular labeling data and/or the consensus data. In some embodiments, the medical scan labeling quality assurance system can aggregate labeling data assigned to a medical scan by each user for display to the expert as a histogram, indicating how many users agreed on each label for a medical scan. This information can be sent to the expert client device for display in conjunction with each medical scan and/or in conjunction with display of raw labeling data from the N users, to further assist the expert in determining the golden labeling data for each medical scan.

The interactive interface **3177** can further allow the expert user to enter correction data for the labeling data submitted by the N users. The correction data can be entered individually for the labeling data of each of the N users, and can be entered individually for each of the M medical scans for each of the N users, where a correction data set for a user corresponds to the correction data for a single user's labeling data for all of the medical scans 1-M. For example, the expert user can enter text in a text box presented by the interactive interface indicating errors in the labeling data for one of the users, allowing the expert user to identify reasons that the labeling data was correct or incorrect, and/or to identify reasons the labeling data entered by the user is associated with common pitfalls. The expert user can interact with interactive interface **3177** to draw on or otherwise superimpose shapes, lines, and text over the medical scan itself, for example, to indicate abnormalities the user missed and/or to indicate benign features that were improperly identified as abnormalities. The correction data can indicate any other problems with the labeling data entered by the user, can indicate feedback to the user that generated the labeling data, and/or can include positive and/or encouraging comments or other feedback indicating the user generated completely correct or mostly correct labeling data. The correction data can further include recommendations for areas that the user should study for further review, for example, indicating fields of medicine, types of medical scans, types of prompts of one or more prompt decision trees, types of anatomical regions and/or types of abnormalities corresponding to weak performance and/or corresponding to areas where the user could use more practice. In this fashion, the correction data can be intended as a learning tool for the user to improve in the future. In some embodiments, aggregate correction data is generated based on the expert user's evaluation of a user's performance overall over the medical scans 1-M. This can be determined based on common pitfalls or errors identified by the expert user based on the set of labeling data for the set of medical scans viewed

as a whole, and/or based on strengths and/or trends that were consistently performed well by the user over the set of medical scans as a whole.

Each set of correction data 1-N can also indicate scoring data identified by the expert user. This can include overall scoring data, scoring data for each of the set of labeling data for a user, scoring points awarded for strong performance, and/or scoring points deducted for errors identified by the expert in the labeling data. The expert user can identify a ranking of the N users' performance for each medical scan or for review over all medical scans, and/or each set of correction data can indicate the ranking identified by the expert relative to the other N-1 users.

As shown in FIG. 13D, the medical scan labeling quality assurance system 3004 can receive the golden labeling data for each of the M medical scans, and can receive correction data sets 1-N indicating correction data sets for each user. A scoring step 3150 can be performed to determine performance score data for each of the labeling data sets 1-N. The performance score data for each of the N users can be generated automatically by the medical scan labeling quality assurance system 3004 based on comparing the set of golden labeling data to the set of labeling data for each of the N users.

The performance score data generated for a user can be an aggregate score for labeling data of all of the medical scans in the set. An individual performance score can be calculated by comparing labeling data for each of the M medical scans for a user to the corresponding golden labeling data of the set of golden labeling data. The performance score data for a user can be an aggregate score based on these M individual performance scores. This can include an average score of the M individual performance scores, can include a median score of the M individual performance scores, can include a maximum score of the M individual performance scores, can include a minimum score of the M individual performance scores, and/or can include some other overall score based on the M individual performance scores. In some embodiments, the performance score data preserves each of the M individual performance scores.

More favorable scores can be generated for labeling data that is the same as, or deviates only slightly from, the corresponding golden labeling data. Less favorable scores can be generated for labeling data that is different from, or deviates greatly from, the corresponding golden labeling data. For example, a difference between the labeling data and corresponding golden labeling data for a single medical scan can be calculated. This can include calculating the Euclidian distance between a feature vector of corresponding to the user's labeling data for a medical scan and the feature vector of the corresponding golden labeling data, where a higher performance score is assigned to a user whose labeling data is a smaller Euclidian distance from the corresponding golden labeling data, and a lower performance score is assigned to a user whose labeling data is a larger Euclidian distance from the golden labeling data.

Alternatively or in addition, the performance score for an individual medical scan can be presented as a binary indicator, indicating whether or not the labeling data was correct based on a comparison to the golden labeling data. The binary indicator can be determined based on absolute correctness. Alternatively, the binary indicator or can indicate whether or not a difference of the labeling data from the golden labeling data was small enough to compare favorably to a correctness threshold, where the binary indicator indicates whether or not the labeling data was "mostly correct" in accordance with the correctness threshold.

The performance score data can further be generated based on correct labeling of each of a plurality of differential diagnoses indicated in golden labeling data for a medical scan. For example, if golden labeling data for one of the medical scans indicates X abnormalities, a corresponding performance score for each user for this medical scan can be generated based on a number or proportion of the X abnormalities that were correctly identified.

The performance score data can further be generated based on the set of correction data corresponding to each set of labeling data. For example, the scoring data and/or rankings determined by the expert upon review of the sets of labeling data can be utilized in generating the performance score data. The performance score data can be set equal to the scoring data and/or rankings set by the expert user. Alternatively, generating the performance score data for a user can include weighing the scores and/or ranks assigned by the expert user with a first weight, and weighing the automatically generated scoring data based on the calculated difference between the set of labeling data of the user and the set of golden labeling data with a second weight. The performance score data is calculated in accordance with summing the product of the first weight and the expert determined scoring data with the product of the second weight and the automatically calculated scoring data.

Alternatively or in addition, the performance score data can further be generated based on the user's performance relative to other users and/or relative to average performance. Consider an example where most users incorrectly labeled a first medical scan, but a first user labeled the first medical scan correctly. This can warrant a more favorable performance score for the first example than a second example where all or almost all users, including the first user, correctly labeled a second medical scan. The first medical scan is determined to be a more challenging medical scan to label correctly than the second medical scan, which is determined to be easy to label correctly. Thus, the first user is rewarded more heavily for correctly labeling the first, challenging medical scan than the second, easy medical scan. The performance score data for the first user includes a first performance score for labeling data of the first medical scan that is more favorable than a second performance score for labeling data of the second medical scan, even when the labeling data for both the first medical scan and second medical scan are equally "correct" relative to the golden labeling data, to reward the first user more heavily for their correct labeling of the more challenging medical scan.

Consider this same example with this first, challenging medical scan that was labeled incorrectly by almost all users and this second, easy medical scan that was labeled incorrectly by almost all users. A second user in the set of users was among the users that incorrectly labeled the first medical scan, and was also the only user to incorrectly label the second medical scan. Incorrectly labeling the easy medical scan can warrant a more unfavorable performance than incorrectly labeling the challenging medical scan. Thus, the second user is penalized more heavily for incorrectly labeling the second, easy medical scan than the first, challenging medical scan. The performance score data for the second user includes a first performance score for labeling data of the first medical scan that is more favorable than a second performance score for labeling data of the second medical scan, even when the labeling data for both the first medical scan and second medical scan are equally "incorrect" relative to the golden labeling data, to penalize the second user more heavily for their incorrect labeling of the easy medical scan.

Following this mentality, some or all of the performance score data for individual medical scans for each user can be generated based on how similar or dissimilar a user's labeling data is from the most popular labeling data, from consensus labeling data, and/or from average labeling data. A histogram generated for labeling data received from the N users for each medical scan can also be utilized to determine how common or uncommon a user's labeling data is. A second difference between the labeling data of a medical scan and corresponding consensus labeling data, corresponding average labeling data, and/or corresponding most common labeling data can be calculated. This can include calculating the Euclidian distance between a feature vector of corresponding to the user's labeling data for a medical scan and the feature vector of the corresponding consensus labeling data, corresponding average labeling data, and/or corresponding most common labeling data.

This second calculated difference can be utilized in conjunction with the first calculated difference from the golden labeling data, and/or can be utilized in conjunction with a binary indicator indicating whether or not the labeling data was correct based on a comparison to the golden labeling data, and/or indicating whether or not a difference of the labeling data from the golden labeling data was small enough to compare favorably to a correctness threshold. As a specific example, consider a neutral score of 0. The performance score can be generated by assigning a score of 1 to correct (or mostly correct) labeling data and by assigning a score of -1 to incorrect (or mostly incorrect) labeling data for a single medical scan. A performance score for the single medical scan can be calculated multiplying D, denoting the magnitude of the calculated second difference from average, consensus, or most common data, by the assigned score of 1 or -1 based on whether this labeling data was correct or incorrect, resulting in D or -D. The performance score data over all of the medical scans M for a single user can be computed by calculating this single performance score for the user's labeling data for each medical scan, and summing the M calculated performance scores.

Some or all of the correction data discussed herein can be generated automatically by the medical scan labeling quality assurance system and/or the client device locally in conjunction with executing application data received from the medical scan labeling quality assurance system that causes the client device to generate the correction data automatically. For example, the performance score data generated for each user can be utilized to generate the set of correction data for each user, and/or a comparison of the labeling data to the golden labeling data for each medical scan of a user can be utilized to generate correction data for each medical scan for each user. Some or all of the automatically generated correction data can be displayed to the expert client device for review as suggested correction data, and the expert client device can utilize this suggested correction data in generating the final correction via the interactive interface **3177**, for example, by providing edits to and/or additions to the suggested correction data.

Once the performance score data is generated for each user, each of the N performance score data is sent to a client device of a respective one of the N users via the network, as shown in FIG. **13D**. Each user can view their respective performance score data indicating their overall performance and/or individual scores 1-M via the interactive interface **3175**. For example, the interactive interface can present the set of medical scans 1-M and the respective individual performance scores to the user. In some embodiments, ranking information indicating how well the user performed

relative to the other N-1 users is also transmitted to the user, allowing the user to gauge their relative performance for each medical scan individually and/or to gauge their overall relative performance. Furthermore, all of the performance score data can be sent to all of the client devices allowing the user to see how much their individual performance scores for each medical scan and/or all performance scores deviate from those of the other N-1 users.

Each set of correction data generated via expert user input and/or generated automatically can also be transmitted to a client device of a respective one of the N users via the network, as shown in FIG. **13D**. The set of correction data can be displayed to the user in conjunction with the user's set of labeling data and/or the set of medical scans. For example, each medical scan can be displayed one at a time to the user in sequence, along with the correction data and/or the user's original labeling data. This can include displaying the golden labeling data in conjunction with the user's labeling data, for example, where the user's labeling data is distinguishable from the golden labeling data and/or where differences between the user's labeling data from the golden labeling data is visually indicated. This can include displaying any text and/or superimposed drawings of the correction data in conjunction with each medical scan.

In some embodiments, data corresponding to other users can be sent to some or all additional users of the set of users 1-N. A client device of a first user can receive sets of labeling data generated by other users, performance score data for other users, and/or set of correction data generated for other users. In some embodiments, all of the sets of labeling data 1-N, all of the performance score data 1-N, and/or all of the set of correction data 1-N are sent to all of the users 1-N. The interactive interface **3175** can display the labeling data, performance score data, and/or correction data of other users in the same fashion that they can view their own labeling data, performance score data, and/or correction data, for example, where this information is displayed in conjunction with display of each medical scan one at a time. In some embodiments, the user can view their own labeling data and labeling data of one or more other users, for example, in the same fashion that the interactive interface **3177** displays labeling data for some or all of the users at the same time, where different labeling data of different users is distinguishable. In some embodiments, the histogram data, consensus labeling data, average labeling data, common labeling data, and/or other aggregate analysis of all of the sets of labeling data 1-N can be viewed by the user, allowing the user to recognize how similar or dissimilar their labeling data is from that of other users.

In some embodiments, the expert user can indicate custom teaching examples for a user based on weaknesses of the user in the correction data for one or more of the M medical scans. The expert user can identify labeling data of at least one different user for the one or more M medical scans as some or all of the custom teaching example data included in the correction data. The medical scan labeling quality assurance system **3004** can send only labeling data identified as teaching examples to the corresponding user based on the indication in the correction data. The medical scan labeling system can also send the corresponding individual performance scores and/or the corresponding correction data for the labeling data of the indicated custom teaching examples.

The expert user can also indicate global teaching examples that would be appropriate for any user. For example, the expert user can identify labeling data for at least one medical scan that succumbs to a common pitfall and/or that avoids a common pitfall. As another example, the

expert user can identify labeling data corresponding to a strong performance that serves as a good example of proper labeling and/or corresponding to a weak performance that serves as an example of poor labeling that should be avoided. The labeling data identified as the global teaching data can be sent to all of the client devices 1-N for display via the interactive interface. In some embodiments, the corresponding correction data and/or individual performance scores of the global teaching examples can also be transmitted to all of the client devices 1-N for display via the interactive interface. In some embodiments, the labeling data, corresponding correction data, and/or corresponding performance score data of the global teaching can be made available to additional users in the system, and can be transmitted to additional corresponding client devices. The labeling data, corresponding correction data, and/or corresponding performance score data can be stored as historical global teaching examples and can be fetched or sent automatically to additional users. For example, the labeling data, corresponding correction data, and/or corresponding performance score data can be mapped to the corresponding medical scan in the medical scan database, and/or can be stored in other memory of the medical scan labeling quality assurance system **3004** and/or in another database of the database storage system **140**.

The medical scan quality assurance labeling system can automatically determine some or all of the custom teaching examples by identifying custom subsets of labeling data of other users based on the performance score data, where the custom subsets are determined to be most appropriate based on the user's errors, based on the types of medical scans the user struggled with, and/or based on the types of abnormalities the user struggled to identify properly. Only the custom subsets of labeling data of other users are transmitted to the corresponding user for display via the interactive interface. The corresponding correction data and/or individual performance scores of the custom subsets of labeling data can also be transmitted to the corresponding user for display via the interactive interface. In some embodiments, the custom subsets are identified in the historical global teaching data and/or other historical labeling data from a previous quality assurance process.

The medical scan labeling quality assurance system **3004** can also automatically determine some or all of the global teaching examples sent to users and/or stored in a database as historical data. For example the medical scan labeling quality assurance system **3004** can automatically determining common pitfalls and selecting labeling data that succumbs to the common pitfall and/or that avoids the common pitfall. For example, the common pitfalls can be identified based on the common labeling data of the histogram data that compares unfavorably to the global labeling data and/or is otherwise determined to be incorrect. Common pitfalls can also be identified by the expert and can be received from the expert client device, and can be utilized to identify common pitfalls in the sets of labeling data. and can be further identified based on determining common labeling data that As another example, the medical scan labeling quality assurance system **3004** can select labeling data corresponding to highest scoring labeling data and/or that corresponds to lowest scoring labeling data for one or more medical scan categories and/or one or more abnormality types as global teaching examples.

FIG. **13E** illustrates transmitting the golden labeling data 1-M to the medical scan database. For example, golden labeling data for the M medical scans and/or golden labeling data for other medical scans generated by experts in other

performances of the quality assurance process can be utilized to generate a golden training set to train a medical scan analysis function, for example as discussed in conjunction with FIG. **12F**. As another example, the medical scan image analysis system **112** can utilize the golden labeling data as output labels for a neural network in training the model as discussed in FIG. **7A**. The trained medical scan analysis function can be utilized by one or more subsystems **101** and/or by the medical picture archive integration system **2600** to generate diagnosis data and/or other inference data for new medical scans.

FIG. **13E** also illustrates transmitting the performance score data for each of the users to the user database. A user profile entry of each user can be updated based on this received user profile data. The performance score data can correspond to accuracy data **531** of the user's performance score data **530**, and/or can correspond to categorized performance score data **534** based on the corresponding category of the medical scan. The performance score data **530** of the user's user profile entry can be updated to reflect the most recent performance score data, can reflect an average of all of the performance score data generated for that user for all of their performance score data generated over time, and/or can otherwise be updated based on the new performance score data received from the medical scan labeling quality assurance system **3004**. As discussed previously, as the performance score data is updated for a user over multiple quality assurance processes, the performance score data of the user can dictate whether the user is identified as an expert, whether the user is included in the set of users for future quality assurance processes, and/or whether the user is put on probation and/or banned from generating labeling data for other medical scans received as part of the normal functioning of the system.

The sets of correction data 1-N can also be sent to the user database or another database, for example, allowing each user to fetch their own set correction data in the future and/or allowing other users to fetch correction data for review. This can further allow historical sets of correction data to be fetched for custom or global teaching examples identified by the expert user and/or identified by the medical scan labeling quality assurance system **3004** automatically for transmission some or all of the N users and/or additional users. In some embodiments, each original set of labeling data is also stored in conjunction with the corresponding set of correction data.

In various embodiments, a medical scan labeling quality assurance system includes a medical scan database that includes a plurality of medical scans, a user database that includes a plurality of user profiles corresponding to a plurality of users of the medical scan labeling quality assurance system, a processing system that includes a processor, and a memory that stores executable instructions. The executable instructions, when executed by the processing system, cause the medical scan labeling quality assurance system to select a set of users from the user database in response to determining a scheduled interval has elapsed. A set of medical scans are selected from the medical scan database. The set of medical scans are transmitted, via a network, to a set of client devices associated with the set of users. The set of medical scans are displayed to the set of users via a first interactive interface displayed by a set of display devices corresponding to the set of client devices. A set of labeling data are received from each of the set of client devices via the network. Each set of labeling data is generated by a corresponding one of the set of client devices, and each set of labeling data includes labeling data for each of

the set of medical scans. The labeling data for each of the set of medical scans is generated by the corresponding one of the set of client devices in response to at least one prompt to provide the labeling data via the first interactive interface in conjunction with display of the each of the set of medicals scan.

An expert user is selected from the user database. The expert user is not included in the set of users. The set of medical scans is transmitted, via the network, to an expert client device associated with the expert user. The set of medical scans are displayed to the expert user via a second interactive interface displayed by an expert display device corresponding to the expert user. A golden set of labeling data is received from the expert client device via the network. The golden set of labeling data is generated by the expert client device and includes golden labeling data for each of the set of medical scans. The golden labeling data for each of the set of medical scans is generated by the expert client device in response to at least one prompt to provide the golden labeling data via the second interactive interface in conjunction with display of the each of the set of medicals scans.

A set of performance score data is generated by generating performance score data for each corresponding set of labeling data by comparing the labeling data of each set of labeling data to the golden labeling data of the set of golden labeling data. Each performance score data of the set of performance score data is assigned to a corresponding one of the set of users that generated the corresponding set of labeling data. Each of a set of user profile entries is updated in the user database for each corresponding one of the set of users based on the performance score data of the set of performance score data assigned to the corresponding one of the set of users. Each performance score data of the set of performance score data is transmitted to a corresponding one of the set of client devices for display, via the first interactive interface, to a corresponding one of the set of users to which the each performance score data is assigned.

FIGS. 14A-14B present embodiments of medical scan annotator system 106, for example, when utilized in conjunction with the medical scan hierarchical labeling system 3002 and/or the medical scan labeling quality assurance system 3004. As illustrated in FIG. 14A, the medical scan annotator system 106 can select a medical scan from the medical scan database 342 for transmission via network 150 to one or more client devices 120 associated with a selected user set 4010 corresponding to one or more users in the user database 344. A medical scan can be selected for annotation based on an assigned priority and/or based on a turn-based queue, for example, based on the scan priority data 427 of the corresponding medical scan entry 352. The client device 120 of each user of the selected user set 4010 can display one or more received medical scans to the via the interactive interface 275 displayed by a display device corresponding to the client device 120, for example, by displaying medical scan image data 410 in conjunction with the medical scan assisted review system 102.

The interactive interface 275 displayed by client devices 120 of each user in the selected user set 4010 can include a prompt to provide annotation data 4020 corresponding to the medical scan. This can include a prompt to provide a text and/or voice description via a keyboard and/or microphone associated with the client device. This can also include a prompt to indicate one or more abnormalities in the medical scan, for example, by clicking on or outlining a region corresponding to each abnormality via a mouse and/or touchscreen. For example, the interactive interface can

prompt the user whether or not an abnormality is present. If the user indicates an abnormality is present, the interactive interface can prompt the user to identify the region that includes the abnormality. This can include allowing the user to scroll through one or more slices, to identify one or more slices that contain the abnormality, and to select a region of the one or more slices that contains the abnormality. Once the region is identified, the interactive interface can prompt the user to provide descriptive information classifying an abnormality based on its size, type, etc. To aid the user in providing this information, the user interface can automatically crop one or more slices based on the identified region and/or zoom in on the identified region. In various embodiments, the medical scan can be presented for annotation by utilizing the medical scan assisted review system 102, for example, presented in the new annotation mode. The interactive interface 275 can present the medical scan by utilizing interface features indicated in the display parameter data 470 and/or the interface preference data 560 of the user, and/or the user can indicate the annotation data via the interactive interface 275 by utilizing interface features indicated in the display parameter data 470 and/or the interface preference data 560 of the user. For example, some or all of the annotation data 4020 can correspond to, or be automatically generated based on, user input to the interactive interface.

Annotation data 4020 can be transmitted from each client device of users in the selected user set 4010 to the medical scan annotator system 106, for example, in response to receiving input data via the interactive interface indicating that the annotations are complete. The annotation data 4020 can be raw annotation data corresponding directly to the user input, or can be further processed by the client device before transmission. For example, a more precise region corresponding to each abnormality can be determined automatically based on the user input and by determining actual boundary points of the abnormality by utilizing image processing techniques and/or text and/or voice input can be processed and/or parsed, for example, by utilizing a medical scan natural language analysis function and/or medical report analysis function to generate medical codes 447 or other diagnosis data 440 corresponding to the medical scan. Such processing can also be performed by the medical scan annotation system 106 and/or another subsystem when the raw annotation data is received.

The medical scan annotator system 106 can evaluate the set annotation data 4020 received from the selected user set 4010 to determine if a consensus is reached, and/or generate a final consensus annotation 4030, for example, by performing an annotation consensus function 4040. For example, consider a selected user set 4010 that includes three users. If two users annotate a medical scan as "normal" and the third user annotates the medical scan as "contains abnormality", the annotation consensus function 4040 performed by medical scan annotator system 106 may determine that the final consensus annotation 4030 is "normal" by following a majority rules strategy. Alternatively, the medical scan annotator system 106 can determine that a consensus is not reached because one of the users indicated that an abnormality is present, and that the medical scan should not be passed off as normal because a level of confidence that the scan is normal, determined by a calculated consensus confidence score 4050, does not exceed a consensus confidence threshold. The confidence thresholds required for consensus can differ for different types of scans and/or severity of diagnosis.

If the medical scan annotator system 106 determines that a consensus is achieved, it can automatically generate the

final consensus annotation **4030**, and can map this final consensus annotation to the medical image in the medical scan database in diagnosis data **440**, and/or transmit the consensus annotation to an originating entity of the medical scan. The medical scan annotator system **106** can also map the calculated consensus confidence score to the medical image in the confidence score data **460**. In some embodiments, a truth flag **461** will automatically be assigned to all final consensus annotation **4030** in the confidence score data **460** and/or will automatically be assigned to final consensus annotation **4030** that exceeds a truth threshold. In some embodiments, annotation data **4020** received from each user and/or a corresponding annotation confidence score can also be stored in the medical database, mapped to the corresponding user and/or the corresponding performance score in the annotation author data **450**.

In some embodiments, for example where annotation data **4020** includes several attributes, the annotation consensus function **4040** performed by the medical scan annotation system **106** can determine whether a consensus is reached by calculating a difference between two or more received annotation data **4020**, for example, by generating a feature vector for annotation data **4020** received from each user. Each feature vector can be generated based on keywords, medical codes, abnormality location in the medical scan, abnormality size and/or shape in the medical scan, a classification of the abnormality, or other attributes listed in annotation data **4020** received from each user. Performing the annotation consensus function **4040** can further include calculating the Euclidian distance or other vector distance between the two or more feature vectors. Performing the annotation consensus function **4040** can further include determining if consensus is reached by determining if the average of these Euclidian distances is below a certain discrepancy threshold, for example, after determining and removing outlier annotations from the set. Similarly, the annotation consensus function **4040** can further include determining if consensus is reached by first generating the final consensus annotation **4030**, and then calculating the Euclidian distance between each annotation feature vector and the final consensus annotation **4030**, where consensus is determined to be reached and the final consensus annotation is confirmed only if the average of these calculated Euclidian distances is below a certain discrepancy threshold. The annotation consensus function **4040** can calculate the final consensus annotation **4030** itself by creating a consensus feature vector, where each attribute of the consensus feature vector is determined by calculating a mean, median or mode of each corresponding annotation feature extracted from all of the received annotation data **4020**. In this fashion, calculating the consensus confidence score **4050** can include calculating such an average Euclidian distance, where distances with larger magnitudes correspond to lower or otherwise less favorable consensus confidence scores **4050**, and where distances with smaller magnitudes correspond to higher or otherwise more favorable consensus confidence scores **4050**. Alternatively or in addition, the final consensus annotation **4030** can be generated based on the most closely matching annotations and/or based on another average, for example, calculating an average identified region that includes an abnormality.

The annotation consensus function **4040** further determine whether or not consensus is reached based on overall or categorized performance score data **530** and/or qualification data **540** of each user in the selected user set **4010**. For example, each annotation data **4020** can be weighted based the performance scores and/or qualifications of the corre-

sponding user. In the example where two users annotate a medical scan as “normal” and a third user annotates a medical scan as “contains abnormality”, the medical scan annotator system **106** may determine that the consensus is “contains abnormality” based on the third user having a much higher performance score and/or being more highly qualified than the first two users. The final consensus annotation **4030** can be generated based on the annotation received from a user with the highest ranking in the category corresponding to the medical scan. The final consensus annotation **4030** can be generated based on calculating a weighted average annotation by computing a weighted consensus feature vector, where feature vectors of higher ranked users receive a higher weight. In some embodiments, each feature of the feature vector can be computed using a different set of user weights, for example, where the different feature weights for each user is determined based on corresponding category-based performance score data and/or qualification data.

Alternatively or in addition, the performance score data associated with the interface features of the interactive interface **275** used by each user to annotate the image can also be utilized to weight the different annotations in reaching consensus. Such weights can be applied when generating a consensus feature vector, where each annotation feature vector is weighted according to the performance score data of one or more corresponding interface features used by the corresponding user.

In some embodiments, confidence scores for each individual annotation can also be calculated for each user’s annotation, and the consensus confidence score **4050** can be generated based on these confidence scores, for example, based on an average confidence score, based on confidence scores of annotation data that matches the final consensus annotation **4030**, etc. In some embodiments, the final consensus annotation **4030** can be generated based on these confidence scores, for example, where annotation feature vectors are weighted based on a corresponding confidence score. The confidence scores for each annotation data **4020** can be generated automatically, for example, based on performance score data **530** as discussed herein. Individual confidence scores and/or a consensus confidence score **4050** can also be updated retroactively as new annotation data is received, for example, if new annotation data is received from another user, for example corresponding to an expert review when consensus is not reached, and/or if new annotation data is automatically generated by a subsystem after the consensus data is generated.

The medical scan annotator system **106** can also utilize auto-generated annotation data of the medical scan to determine if consensus is reached and/or to generate the final consensus annotation **4030**. The auto-generated annotation data can be automatically generated by medical scan annotator system **106** by utilizing one or more medical scan analysis functions. The auto-generated annotation data can also be retrieved from the medical scan database **342** if it was generated by a subsystem **101** previously. One or more auto-generated annotations can be assigned their own weights and/or confidence scores, for example, based on the model accuracy data **631** and/or another determined performance of the function and/or subsystem responsible for creating each auto-generated annotation. Each auto-generated annotation data can be thus treated as an annotation from another user, and can be used to determine if consensus is reached and/or to generate the consensus annotation in the same fashion.

Alternatively, the auto-generated annotation can be merely verified based on the annotation data **4020** received from the selected user set **4010** by determining that the user annotations are close enough to the auto-generated annotation based on the discrepancy threshold. For example, this process may be utilized by the medical scan diagnosing system **108** to perform the output quality assurance step. The auto-generated annotation can be sent to the selected user set **4010** as part of this verification process, for example, displayed by each interactive interface **275** in conjunction with the medical scan assisted review system **102** as displayed annotation data, and the annotation data **4020** received from the selected user set **4010** can be include verification of and/or corrections of the auto-generated annotation. Alternatively, the medical scan can be sent without the auto-generated annotation and/or the auto-generated annotation can be hidden from view as part of a blind review, to ensure that the users are not biased in creating annotation data by the auto-generated annotation.

FIG. **14B** illustrates an embodiment of the medical scan annotator system **106** upon determining that a consensus is not achieved, for example, because the calculated consensus confidence score **4050** does not exceed the consensus confidence threshold. The medical scan annotator system can select an expert user, for example, a user whose qualification data **540** indicates they are an expert in the category corresponding to the medical scan or who otherwise is identified as an expert based on their performance score data. The expert can receive the medical scan on a corresponding client device and annotate the image, for example, where the interactive interface **275** displays the medical scan image data **410** in conjunction with the medical scan assisted review system **102**, and where the interactive interface utilizes interface features indicated in the display parameter data **470** of the medical scan and/or indicated in the interface preference data **560** of the user profile entry **354** of the expert user. The expert can view the annotation data **4020** generated by the selected user set **4010**, for example, presented as the displayed annotation data of the medical scan assisted review system **102**. Annotation data **4020** of each user can be displayed one at a time and the expert user can elect to advance to the next user's annotation data **4020**. Alternatively, all of the annotation data **4020** can be displayed simultaneously for example, in different colors corresponding to each user's annotations and/or overlaid as translucent, highlighted regions, for example, where a portion of the highlighted region is more opaque when multiple users agree that the portion is included in the abnormality. In other embodiments, the annotation data **4020** can be hidden from the expert user, and the expert user can enter their own annotations in conjunction with a blind review to reduce bias.

Expert annotation data **4070** can be generated automatically, and can be transmitted automatically to the medical scan annotation system **106**. The medical scan annotator system can automatically assign the received expert annotation data **4070** as the final consensus annotation **4030**, and/or can assign a truth flag **461** to the expert annotation data **4070** in the confidence score data **460** of the medical scan. Alternatively, the expert annotation data **4070** can be compared to the previous annotation data **4020** and determine if consensus has been reached. For example, the expert annotation data **4070** and the annotation data **4020** can be collectively utilized by the annotation consensus function **4040**, where the expert annotation data **4070** is assigned its own, higher weight than the other annotations. If consensus has still not been reached, the medical scan annotation

system can continue to transmit the image other users and processing received annotations until consensus is reached, for example, selecting a new selected user set **4010** and/or selecting a new expert user.

The user profile entries **354** of each user in the selected user set **4010** and/or each expert user can be automatically updated by the medical scan annotator system **106** or another subsystem **101** by generating and/or updating performance score data **530** for each user based comparing their annotation to the final consensus annotation **4030**. For example, the accuracy score data **531** of the performance score data **530** can be generated by calculating the Euclidian distance between a feature vector of a user's annotation and the feature vector of the consensus annotation as described previously, where a higher performance score is assigned to a user whose annotation is a smaller Euclidian distance from the consensus, and a lower performance score is assigned to a user whose annotation is a larger Euclidian distance from the consensus. The efficiency score data **532** of the performance score data can be automatically generated, for example, based on an annotation duration determined based on a difference between a first time that each user received the medical scan and a second time each user completed the annotation. The efficiency score data **532** can be further based on a difference between the annotation duration of each user and an average annotation duration computed for annotation durations of the selected user set. Aggregate performance data for each user can be generate and/or updated based on past accuracy and/or efficiency scores, based on how many scans have been annotated in total, based on measured improvement of the user over time, etc. Similarly, the performance score data **630** corresponding to medical scan analysis functions utilized to generate the auto-generated annotation data can be generated and/or updated by comparing the auto-generated annotation data to the final consensus annotation **4030** in a similar fashion and/or by comparing the computed annotation duration of a corresponding medical scan analysis functions to other computed annotation durations of other medical scan analysis functions that generated auto-generated annotation data for the medical scan.

The selected user set **4010** can be selected based on the performance score data **530** and/or qualification data **540** of each user corresponding to previous uses only the medical scan annotation system **106**, or corresponding to usage of several subsystems **101**. For example, a medical professional with a user profile indicating that he/she ranks above a certain threshold in annotating CT scans and/or indicating that he/she is highly qualified in the study of the lungs can be automatically selected by the medical scan annotator system to annotate a triaged medical scan identified as a lung CT scan. The size of the selected user set **4010** that receive a medical scan can be optimized based on the quality of the users selected, for example, based on calculating the probability of reaching consensus and/or calculating the probability that a consensus confidence score will be above a confidence threshold, and ensuring the probability falls above a probability threshold. For example, a first medical scan can be sent to a two medical professionals with high scores, qualifications, rankings, or correct annotation percentages. A second medical scan may be sent to ten medical professionals with lower scores or qualifications based on calculating that the probability of a correct consensus probability falls above a probability threshold.

In some embodiments, the medical scan annotator system **106** can first select a medical scan for annotation automatically, and in response, the selected user set **4010** can be

determined automatically to annotate the selected medical scan based on determining users with highly ranked overall scores and/or based on categorized performance data 534 and/or qualification data 540 that corresponds to an identified scan classifier data 420 of the selected medical scan. Alternatively or in addition, the selected user set 4010 can be determined based on the size of a queue of medical scans already assigned to each user. For example, the selected user set 4010 can correspond to users with matching qualifications that correspond to the scan classifier data 420 and/or correspond to users with the lowest queues of other medical scans to annotate.

In other embodiments, the medical scan annotator system 106 can first determine one or more available users automatically, for example, based on medical scan queue lengths for each user in the system and/or in response to one or more users requesting to annotate a medical scan. In such cases, some or all of these identified users can be added to the selected user set 4010, and the medical scan can be selected based on corresponding categorized performance data 534, qualification data 540 or other relevant user profile data of users in the selected user set 4010.

FIGS. 14C-14V present example embodiments of a user interface of a medical scan annotator system 106, for example, presented in conjunction with the medical scan assisted review system 102. Some or all features presented in FIGS. 14C-14V can also be utilized in conjunction with other subsystems and can be included in the interface features. FIGS. 14C-14G present interface features for chest CT nodule characterization, and can be displayed in conjunction with a chest CT scan. Annotation data 4020 can be generated based on user selections in the user interface, and can be used to populate abnormality classification data 445 for abnormality classifier categories 444 such as “nodule spiculation”, “nodule lobulation”, “nodule texture”, “nodule calcification”, “nodule sphericity” and/or “nodule internal structure” for the associated medical scan. FIGS. 14H-14J present interface features for presentation to a user in conjunction with an identifying chest CT nodule, allowing a user to add new contours for one or more scans for a patient, for example, over multiple years, and indicate malignancy. As shown in FIG. 14K, the scan can be presented in conjunction with these interface features. FIGS. 14L-14O present interface features for presentation to a user in conjunction with identifying abnormalities in a chest x-ray. Users can classify each abnormality and draw a shape around each abnormality in the scan.

FIG. 14P presents a view of a chest x-ray presented via the interface before a user identifies regions of interest, and FIG. 14Q presents a view of the chest x-ray via the interface after the user identifies regions of interest of multiple abnormalities, indicated by seven polygons 1022. FIG. 14R presents interface features for comparing chest x-ray severity for multiple patients, displayed in conjunction with multiple x-rays that can be displayed in adjacent views or can be displayed one at a time where the user can toggle between them. A user can compare multiple scans corresponding to multiple patients, and provide feedback indicating differences between the patients, comparing if one patient’s case is more severe than another, or determine which of two scans appears to be more normal.

FIGS. 14S-14V present interface features for chest x-ray triage classification, displayed in conjunction with a chest x-ray. A user can select abnormality classification data that can be used to generate annotation data 4020 and/or to populate abnormality classification data 445. As shown, some or all abnormality classification categories displayed,

which can be determined based on abnormality classifier categories 444, can be presented, and hierarchal subcategories can be presented in response to a user selecting one of a plurality of abnormality classification categories that are present.

In some embodiments, the medical scan hierarchical labeling system 3002 is integrated within and/or utilizes features described in conjunction with the medical scan annotator system 106. In particular, the interfaces presented in some or all of FIGS. 14C-14V can be utilized by the medical scan hierarchical labeling system 3002, where options presented in in some or all of FIGS. 14C-14V can correspond to a set of selection options of a node in accordance with a prompt decision tree and/or where prompt decision trees of the medical scan hierarchical labeling system 3002 utilize at least one of the prompts and/or sets of options presented in in some or all of FIGS. 14C-14V as one or more prompts of one or more prompt decision trees, where the annotation data 4040 corresponding to the labeling data generated by the client device. The fixed set of diagnosis, characterization, and/or localization options can include some or all of the selections presented in in some or all of FIGS. 14C-14V, where fields of a medical scan entry correspond to some or all of the selections presented in in some or all of FIGS. 14C-14V and/or have valid entries corresponding to some or all of the selections presented in some or all of FIGS. 14C-14V.

The interactive interface can present hierarchical sets of options as a user advances through a prompt decision tree as presented in FIGS. 14S-14V, where each set of options is presented as an indented list in accordance with advancing to a deeper layer of the prompt decision tree, and where each set of options is not presented until the corresponding selection is made. For a hierarchical decision tree corresponding to the prompts presented in FIGS. 14S-14V, labeling data generated in response to a user selecting “submit and next” at FIG. 14V can indicate a leaf node of a first abnormality, with diagnosis, characterization, and localization data fully described as “pulmonary vasculature”, “plethora”, “diffuse”, “left”, and “lobe-left upper”, and can further indicate a leaf node of a second abnormality with diagnosis, characterization, and localization data fully described as “mediastinum”, “compression of structure”, and “inferior-anterior”.

In some embodiments, the medical scan labeling quality assurance system 3004 is integrated within and/or utilizes features described in conjunction with the medical scan annotator system 106. In particular, the interactive interface 275 of the medical scan annotator system 106 can correspond to the interactive interface 3075 and/or 3175. The annotation data 4040 can correspond to labeling data generated by a client device 120 in conjunction with execution of a labeling application in conjunction with the medical scan hierarchical labeling system 3002 and/or the medical scan labeling quality assurance system 3004. The expert annotation data 4070 medical scan annotator system 106 can correspond to the golden labeling data generated by expert client device 3120. The medical scan labeling quality assurance system 3004 can determine the difference between labeling data and golden labeling data and/or can generate performance score data as described in performing the annotation consensus function 4040 of the medical scan annotator system 106.

FIG. 15A presents a flowchart illustrating a method for execution by a medical scan hierarchical labeling system 3002 that stores executional instructions that, when executed

by at least one processor, cause the medical scan hierarchical labeling system **3002** to perform the steps below.

Step **5002** includes transmitting, via a network, labeling application data to a client device for storage. The labeling data includes a plurality of prompt decision trees. The plurality of prompt decision trees includes a diagnosis prompt decision tree, a characterization prompt decision tree, and a localization prompt decision tree. Each of the plurality of prompt decision trees includes a root node, a set of internal nodes, and a set of leaf nodes. Each root node and each of the set of internal nodes correspond to one of a plurality of prompts. Each root node and each of the set of internal nodes include a set of branches that each correspond to one of a discrete set of selection options for the one of the plurality of prompts. The labeling application further includes application operational instructions that, when executed by at least one client device processor of the client device, cause the client device to execute a labeling application.

Step **5004** includes transmitting, via the network, a medical scan to the client device. Execution of the labeling application by the client device causes the client device to, in response to receiving the medical scan, perform the steps of FIG. **15B**. Step **5006** includes receiving, via the network, a set of labels from the client device, where the set of labels were generated by the client device as a result of the client device performing the steps of FIG. **15B** in accordance with execution of the labeling application. Step **5008** includes populating a medical scan entry of the medical scan in a medical scan relational database based on the set of labels.

FIG. **15B** presents a flowchart illustrating a method for execution by a client device **120** that stores labeling application data received from the medical scan hierarchical labeling system **3002**. The labeling data includes the plurality of prompt decision trees and the application operational instructions that, when executed by at least one processor of the client device **120**, cause the client device **120** to perform the steps below.

Step **5102** includes receiving a medical scan via a network. Step **5104** includes displaying, via an interactive interface presented on a display device associated with the client device for display to a user associated with the client device, image data of the medical scan. Step **5106** includes automatically determining a starting diagnosis prompt by selecting one of the set of internal nodes of the diagnosis prompt decision tree based on an anatomical region of the medical scan and further based on a modality of the medical scan. Step **5108** includes displaying, via the interactive interface, a plurality of diagnosis prompts of the diagnosis prompt decision tree, in succession, beginning with the starting diagnosis prompt, in accordance with corresponding nodes of the diagnosis prompt decision tree until a first one of the set of leaf nodes of the diagnosis prompt decision tree is ultimately selected. The interactive interface progresses to each next one of the plurality of diagnosis prompts by selecting one of the set of branches of each corresponding node in accordance with each of a plurality of corresponding user diagnosis selections, received via user input. Each of the plurality of corresponding user diagnosis selections corresponds to one of the discrete set of selection options for each one of the plurality of diagnosis prompts displayed via the interactive interface.

Step **5110** includes automatically determining a starting characterization prompt by selecting one of the set of internal nodes of the characterization prompt decision tree based on the anatomical region of the medical scan, based on the modality of the medical scan, and further based on the

first one of the set of leaf nodes of the diagnosis prompt decision tree. Step **5112** includes displaying, via the interactive interface, a plurality of characterization prompts, in succession, beginning with the starting characterization prompt, in accordance with corresponding nodes of the characterization prompt decision tree until a first one of the set of leaf nodes of the characterization prompt decision tree is ultimately selected. The interactive interface progresses to each next one of the plurality of characterization prompts by selecting one of the set of branches of each corresponding node in accordance with each of a plurality of corresponding user characterization selections, received via user input. Each of the plurality of corresponding user characterization selections corresponds to one of the discrete set of selection options for each one of the plurality of characterization prompts displayed via the interactive interface.

Step **5114** includes automatically determining a starting localization prompt by selecting one of the set of internal nodes of the localization prompt decision tree based on the anatomical region of the medical scan, and further based on the modality of the medical scan. Step **5116** includes displaying via the interactive interface, a plurality of localization prompts, in succession, beginning with the starting localization prompt, in accordance with corresponding nodes of the localization prompt decision tree until a first one of the set of leaf nodes of the localization prompt decision tree is ultimately selected. The interactive interface progresses to each next one of the plurality of localization prompts by selecting one of the set of branches of each corresponding node in accordance with a plurality of corresponding user localization selections, received via user input. Each of the plurality of corresponding user localization selections corresponds to one of the discrete set of selection options for each one of the plurality of localization prompts displayed via the interactive interface. Step **5118** includes transmitting, via the network, labeling data that includes a set of labels indicating the first one of the set of leaf nodes of the diagnosis prompt decision tree, the first one of the set of leaf nodes of the characterization prompt decision tree, and the first one of the set of leaf nodes of the localization prompt decision tree.

In various embodiments, the labeling application data is transmitted by the medical scan hierarchical labeling system **3002** to plurality of client devices. A plurality of medical scans are also transmitted to the plurality of client devices by the medical scan hierarchical labeling system **3002**. A plurality of sets of leaf node identifiers for the plurality of medical scans are received by the medical scan hierarchical labeling system **3002** from the plurality of client devices via the network. A plurality of medical scan entries of the plurality of medical scans are populated in the medical scan relational database by the medical scan hierarchical labeling system **3002**. Each of the plurality of medical scan entries is populated based on a corresponding set of leaf node identifiers of the plurality of sets of leaf node identifiers.

In various embodiments, a computer vision model is generated by the medical scan hierarchical labeling system **3002** by performing a training step on the plurality of medical scan entries. Output labels of the plurality of medical scan entries utilize structured fields of the medical scan relational database to indicate corresponding sets of leaf node identifiers. An inference function is performed by the medical scan hierarchical labeling system **3002** on a new medical scan by utilizing the computer vision model to generate inference data. The inference data indicates at least one leaf node of the set of leaf nodes of the diagnosis prompt decision tree, at least one leaf node of the set of leaf nodes

of the characterization prompt decision tree, and at least one leaf node of the set of leaf nodes of the localization prompt decision tree. The inference data is transmitted by the medical scan hierarchical labeling system 3002, via the network, to a second client device for display via a second display device of the second client device.

In various embodiments, execution of the labeling application by the client device further causes the client device to display, via the interactive interface, a differential diagnosis prompt to indicate whether differential diagnosis is present in response to selection of the first one of the set of leaf nodes of the diagnosis prompt decision tree. The client device receives a user selection to enter a differential diagnosis in response to the differential diagnosis prompt. In response to the user selection to enter the differential diagnosis, the client device displays, via the interactive interface, a second plurality of diagnosis prompts of the diagnosis prompt decision tree, in succession, beginning with the starting diagnosis prompt, in accordance with corresponding nodes of the diagnosis prompt decision tree until a second one of the set of leaf nodes of the diagnosis prompt decision tree is ultimately selected, where the second one of the set of leaf nodes is different from the first one of the set of leaf nodes. The client device automatically determines a second starting characterization prompt by selecting one of the set of internal nodes of the diagnosis prompt decision tree based on the anatomical region of the medical scan, based on the modality of the medical scan, and further based on the second one of the set of leaf nodes of the diagnosis prompt decision tree. The client device displays, via the interactive interface, a second plurality of characterization prompts, in succession, beginning with the second starting characterization prompt, in accordance with corresponding nodes of the characterization prompt decision tree until a second one of the set of leaf nodes of the characterization prompt decision tree is ultimately selected. The client device displays via the interactive interface, a second plurality of localization prompts, in succession, beginning with the starting localization prompt, in accordance with corresponding nodes of the localization prompt decision tree until a second one of the set of leaf nodes of the localization prompt decision tree is ultimately selected. The set of labels further indicates the second one of the set of leaf nodes of the diagnosis prompt decision tree, the second one of the set of leaf nodes of the characterization prompt decision tree, and the second one of the set of leaf nodes of the localization prompt decision tree.

In various embodiments, execution of the labeling application by the client device further causes the client device to display, via the interactive interface, the differential diagnosis prompt to indicate whether differential diagnosis is necessary in response to selection of each one of a selected plurality of leaf nodes of the set of leaf nodes of the diagnosis prompt decision tree, until receiving a user selection indicating further differential diagnosis is not present. The client device automatically determines a plurality of starting characterization prompts corresponding to the selected plurality of leaf nodes of the set of leaf nodes of the diagnosis prompt decision tree. The client device displays, via the interactive interface, a plurality characterization prompts for each of the selected plurality of leaf nodes of the set of leaf nodes of the diagnosis prompt decision tree, in succession, beginning with each corresponding one of the plurality of starting characterization prompts, in accordance with corresponding nodes of the characterization prompt decision tree until each of a corresponding selected plurality of the set of leaf nodes of the characterization prompt decision tree are ultimately selected. The client device

displays, via the interactive interface, a plurality of localization prompts for each of the selected plurality of leaf nodes of the set of leaf nodes of the diagnosis prompt decision tree, in succession, beginning with the starting localization prompt, in accordance with corresponding nodes of the characterization prompt decision tree until each of a corresponding selected plurality of the set of leaf nodes of the localization prompt decision tree are ultimately selected. The set of labels further indicates the selected plurality of the set of leaf nodes of the diagnosis prompt decision tree, the corresponding selected plurality of the set of leaf nodes of the characterization prompt decision tree, and the corresponding selected plurality of the set of leaf nodes of the localization prompt decision tree.

In various embodiments, execution of the labeling application by the client device further causes the client device to display, via the interactive interface, an abnormality present prompt to indicate whether an abnormality is present in response to receiving the medical scan. A user selection indicating that no abnormality is present is received by the client device in response to the abnormality present prompt. The client device determines to forego display of the plurality of diagnosis prompts, the plurality of characterization prompts, and the plurality of localization prompts in response to the user selection indicating that no abnormality is present. The set of labels indicates that the medical scan is normal in response to the user selection indicating that no abnormality is present.

In various embodiments, one of the plurality of characterization prompts includes a prompt to the user to indicate whether at least one secondary finding is present. In response to the client device receiving a user selection indicating at least one secondary finding is present, successive ones of the plurality of characterization prompts presented by the client device correspond to characterization of the at least one secondary finding. The labeling data indicates at least one leaf node indicating the at least one secondary finding of a primary diagnosis indicated by the first one of the set of leaf nodes of the diagnosis prompt decision tree.

In various embodiments, execution of the labeling application by the client device further causes the client device to, in response to receiving a user selection indicating at least one secondary findings is present, automatically determine a second starting diagnosis prompt by selecting one of the set of internal nodes of the diagnosis prompt decision tree based on the first one of the set of leaf nodes of the diagnosis prompt decision tree. The client device displays, via the interactive interface, a plurality of diagnosis prompts of the diagnosis prompt decision tree, in succession, beginning with the starting diagnosis prompt, in accordance with corresponding nodes of the diagnosis prompt decision tree until a second one of the set of leaf nodes of the diagnosis prompt decision tree is ultimately selected, where the second one of the set of leaf nodes of the diagnosis prompt decision tree is different from the first one of the set of leaf nodes of the diagnosis prompt decision tree. The client device automatically determines a second starting characterization prompt by selecting one of the set of internal nodes of the diagnosis prompt decision tree based on the anatomical region of the medical scan, based on the modality of the medical scan, and further based on the second one of the set of leaf nodes of the diagnosis prompt decision tree. The client device displays, via the interactive interface, a second plurality of characterization prompts, in succession, beginning with the second starting characterization prompt, in accordance with corresponding nodes of the characterization

prompt decision tree until a second one of the set of leaf nodes of the characterization prompt decision tree is ultimately selected. The client device displays, via the interactive interface, a second plurality of localization prompts, in succession, beginning with the starting localization prompt, in accordance with corresponding nodes of the localization prompt decision tree until a second one of the set of leaf nodes of the localization prompt decision tree is ultimately selected. The set of labels further indicates the second one of the set of leaf nodes of the diagnosis prompt decision tree, the second one of the set of leaf nodes of the characterization prompt decision tree, and the second one of the set of leaf nodes of the localization prompt decision tree. The labeling data indicates the first one of the set of leaf nodes of the diagnosis prompt decision tree corresponds to a primary finding. The labeling data further indicates the second one of the set of leaf nodes of the diagnosis prompt decision tree, the second one of the set of leaf nodes of the characterization prompt decision tree, and the second one of the set of leaf nodes of the localization prompt decision tree correspond to a secondary finding of the primary finding. In various embodiments, the primary diagnosis indicated by the first one of the set of leaf nodes of the diagnosis prompt decision tree corresponds to a brain tumor, and the at least one secondary finding indicated by the at least one leaf node correspond to a fracture and/or a brain bleed.

In various embodiments, a subset of the set of internal nodes of the localization prompt decision tree has a plurality of mandatory branches. One of the plurality of localization prompts displayed via the interactive interface corresponds to one of the set of internal nodes included in the subset. Execution of the labeling application by the client device further causes the client device to display ones of the plurality of localization prompts corresponding to nodes branching from each of the plurality of mandatory branches of the one of the set of internal nodes included in the subset until a plurality of leaf nodes of the set of leaf nodes of the localization prompt decision tree are each reached in succession. The set of labels includes the plurality of leaf nodes of the set of leaf nodes of the localization prompt decision tree. In various embodiments, the medical scan is a head CT scan, and the plurality of leaf nodes of the set of leaf nodes indicates at least one lobe and at least one compartment. In various embodiments, the plurality of leaf nodes of the set of leaf nodes indicates exactly one lobe and exactly one compartment.

In various embodiments, automatically determining the starting characterization prompt is further based on the first one of the set of leaf nodes of the diagnosis prompt decision tree. Automatically determining the starting localization prompt is further based on the first one of the set of leaf nodes of the diagnosis prompt decision tree.

In various embodiments, the medical scan includes a plurality of image slices. Execution of the labeling application by the client device further causes the client device to display, via the interactive interface, a slice selection prompt. A user selection that indicates a selected subset of the plurality of images slices of the medical scan is received by the client device. The starting localization prompt is further selected based on the selected subset of the plurality of image slices. In various embodiments, automatically determining the starting diagnosis prompt is further based on the selected subset of the plurality of image slices.

In various embodiments, execution of the labeling application by the client device further causes the client device to display, via the interactive interface, an urgency prompt. A user selection that indicates a urgency ranking of the medical

scan is received by the client device in response to the urgency prompt. The set of labels further indicates the urgency ranking. In various embodiments, the labeling application data further includes a plurality of additional question sets. Execution of the labeling application by the client device further causes the client device to automatically determine an additional question set from the plurality of additional question sets based on based on the anatomical region of the medical scan and further based on the modality of the medical scan.

In various embodiments, execution of the labeling application by the client device causes the client device to, in response to receiving the medical scan generate a plurality of partitioned portions of the medical scan in accordance with a corresponding plurality of anatomical regions by performing a pre-processing step. The client device determines a starting diagnosis prompt, a starting characterization prompt, and a starting localization prompt for each of the plurality of partitioned portions based on one of the corresponding plurality of anatomical regions. Each one of the plurality of partitioned portions of the medical scan are displayed by the client device in succession. For each displayed one of the plurality of portioned portions of the medical scan, a plurality of diagnosis prompts are displayed by the client device in conjunction with the displayed one of the plurality of partitioned portions of the medical scan, in succession, beginning with a corresponding one of the plurality of starting diagnosis prompts, until one of the set of leaf nodes of the diagnosis prompt decision tree is ultimately selected for the displayed one of the plurality of partitioned portions of the medical scan. For each displayed one of the plurality of portioned portions of the medical scan, the client device displays a plurality of characterization prompts in conjunction with the displayed one of the plurality of partitioned portions of the medical scan, in succession, beginning with a corresponding one of the plurality of starting characterization prompts, until one of the set of leaf nodes of the characterization prompt decision tree is ultimately selected for the displayed one of the plurality of partitioned portions of the medical scan. For each displayed one of the plurality of portioned portions of the medical scan, the client device displays a plurality of localization prompts in conjunction with the displayed one of the plurality of partitioned portions of the medical scan, in succession, beginning with a corresponding one of the plurality of starting localization prompts, until one of the set of leaf nodes of the localization prompt decision tree is ultimately selected for the displayed one of the plurality of partitioned portions of the medical scan. The set of labels indicates the one of the set of leaf nodes of the diagnosis prompt decision tree for each displayed one of the plurality of portioned portions of the medical scan, the first one of the set of leaf nodes of the characterization prompt decision tree for each displayed one of the plurality of portioned portions of the medical scan, and the first one of the set of leaf nodes of the localization prompt decision tree for each displayed one of the plurality of portioned portions of the medical scan.

Some or all of the steps performed by the client device **120** can instead, or additionally, be performed by the medical scan hierarchical labeling system **3002**, for example, in response to transmissions from the client device that indicate user input to the interactive interface. Results of some or all of these steps are transmitted back to the client device, and the interactive interface can display some or all of the prompts and/or other information to the user based on these results received from the medical scan hierarchical labeling system **3002**.

FIG. 16 presents a flowchart illustrating a method for execution by a medical scan labeling quality assurance system 3004 that stores executional instructions that, when executed by at least one processor, cause the medical scan labeling quality assurance system 3004 to perform the steps below.

Step 5202 includes selecting a set of users from a user database in response to determining a scheduled interval has elapsed. The user database includes a plurality of user profiles corresponding to a plurality of users, and the set of users is selected from the plurality of users. Step 5204 includes selecting a set of medical scans from a medical scan database. The medical scan database includes a plurality of medical scans, and the set of medical scans is selected from the plurality of medical scans. Step 5206 includes transmitting the set of medical scans, via a network, to a set of client devices associated with the set of users. The set of medical scans are displayed to the set of users via a first interactive interface displayed by a set of display devices corresponding to the set of client devices. Step 5208 includes receiving a set of labeling data from each of the set of client devices via the network. Each set of labeling data is generated by a corresponding one of the set of client devices, and each set of labeling data includes labeling data for each of the set of medical scans. The labeling data for each of the set of medical scans is generated by the corresponding one of the set of client devices in response to at least one prompt to provide the labeling data via the first interactive interface in conjunction with display of the each of the set of medical scans.

Step 5210 includes selecting an expert user from the user database, where the expert user is not included in the set of users. Step 5212 includes transmitting the set of medical scans, via the network, to an expert client device associated with the expert user. The set of medical scans are displayed to the expert user via a second interactive interface displayed by an expert display device corresponding to the expert user. Step 5214 includes receiving a golden set of labeling data from the expert client device via the network. The golden set of labeling data is generated by the expert client device and includes golden labeling data for each of the set of medical scans. The golden labeling data for each of the set of medical scans is generated by the expert client device in response to at least one prompt to provide the golden labeling data via the second interactive interface in conjunction with display of the each of the set of medical scans.

Step 5216 includes generating a set of performance score data by generating performance score data for each corresponding set of labeling data by comparing the labeling data of each set of labeling data to the golden labeling data of the set of golden labeling data. Step 5218 includes assigning each performance score data of the set of performance score data to a corresponding one of the set of users that generated the corresponding set of labeling data. Step 5220 includes updating each of a set of user profile entries in the user database for each corresponding one of the set of users based on the performance score data of the set of performance score data assigned to the corresponding one of the set of users. Step 5222 includes transmitting each performance score data of the set of performance score data to a corresponding one of the set of client devices for display, via the first interactive interface, to a corresponding one of the set of users to which the each performance score data is assigned.

In various embodiments, the medical scan labeling quality assurance system 3004 transmits each set of labeling data to the expert client device. The set of labeling data is displayed

to the displayed to the expert user via the second interactive interface in conjunction with display of the set of medical scans. In various embodiments, medical scan labeling quality assurance system 3004 receives a set of correction data from each of the set of client devices. Each set of correction data corresponds to one set of labeling data, and each correction data of each set of correction data corresponds to one of the set of medical scans. Each set of correction data is generated by the expert client device in response to at least one additional prompt to provide the each of the set of correction data via the second interactive interface in conjunction with display of each corresponding set of labeling data. The set of performance score data is further generated based on the set of correction data. The medical scan labeling quality assurance system 3004 transmits each correction data of the set of correction data to a corresponding one of the set of client devices for display, via the first interactive interface, in conjunction with set of medical scans and the each performance score data.

In various embodiments, each set of correction data includes comment data, and the comment data is generated by the expert client device in response to text entered by the expert user in at least one text box displayed by the second interactive interface. In various embodiments, medical scan labeling quality assurance system 3004 transmits each correction data of the set of correction data to all of the set of client devices for display in conjunction with the set of medical scans. In various embodiments, the medical scan labeling quality assurance system 3004 selects a second set of users from the user database. A set difference between the second set of users and the set of users and is non-null, where at least one of the second set of users is not included in the first set of users. Alternatively or in addition, an intersection between the set of users and the second set of users is non-null. The set of medical scans and the set of correction data is sent to each one of a second set of client devices corresponding to the second set of users for display, via a third interactive interface, in conjunction with set of medical scans.

In various embodiments, the medical scan labeling quality assurance system 3004 determines a common error in the sets of labeling data for at least one of the set of medical scans based on the set of golden labeling data. The medical scan labeling quality assurance system 3004 identifies at least one set of labeling data that includes the common error. Only the at least one of the set of labeling data that includes the common error and only the at least one of the of the set of medical scans corresponding to the common error are transmitted to the set of medical scans, where the third interactive interface indicates the common error. In various embodiments, the medical scan labeling quality assurance system 3004 generates labeling commonality data by comparing the sets of labeling data to each other to determine similar ones of the sets of labeling data for at least one of the set of medical scans. The medical scan labeling quality assurance system 3004 generates common error data based on comparing the labeling commonality data to the golden labeling data to determine a proper subset of the similar ones of the sets of labeling data that compares unfavorably to the golden labeling data. The common error is determined based on the common error data. In various embodiments, the common error is identified by the expert client device in response to at least one additional prompt to provide the common error data via the second interactive interface.

In various embodiments the medical scan labeling quality assurance system 3004 determines that the scheduled interval has elapsed since transmission of the set of medical scans

to the set of client devices. The medical scan labeling quality assurance system **3004** selects a second set of users from the user database after the scheduled interval has elapsed since selection of the set of users in response to determining the scheduled time interval has elapsed. The medical scan labeling quality assurance system **3004** selects a second set of medical scans from the medical scan database in response to determining the scheduled time interval has elapsed. An intersection between the set of medical scans and the second set of medical scans is null. The medical scan labeling quality assurance system **3004** transmits the second set of medical scans, via a network, to a second set of client devices associated with the second set of users. The second set of medical scans are displayed to the second set of users via the first interactive interface displayed by a second set of display devices corresponding to the second set of client devices.

A second set of labeling data is received from each of the second set of client devices via the network. Each second set of labeling data is generated by a corresponding one of the second set of client devices, and each second set of labeling data includes second labeling data for each of the second set of medical scans. The second labeling data for each of the second set of medical scans is generated by the corresponding one of the second set of client devices in response to at least one prompt to provide the second labeling data via the first interactive interface in conjunction with display of the each of the second set of medical scans. The medical scan labeling quality assurance system **3004** selects a second expert user from the user database, where the second expert user is not included in the second set of users. The medical scan labeling quality assurance system **3004** transmits the second set of medical scans, via the network, to a second expert client device associated with the second expert user. The second set of medical scans are displayed to the second expert user via a second interactive interface displayed by a second expert display device corresponding to the second expert user.

The medical scan labeling quality assurance system **3004** receives a second golden set of labeling data from the second expert client device via the network. The second golden set of labeling data is generated by the second expert client device and includes second golden labeling data for each of the second set of medical scans. The second golden labeling data for each of the second set of medical scans is generated by the second expert client device in response to at least one prompt to provide the second golden labeling data via the second interactive interface in conjunction with display of the each of the second set of medical scans.

The medical scan labeling quality assurance system **3004** generates a second set of performance score data by generating second performance score data for each corresponding second set of labeling data by comparing the second labeling data of each second set of labeling data to the second golden labeling data of the second set of global labeling data. The medical scan labeling quality assurance system **3004** assigns each second performance score data of the second set of performance score data to a corresponding one of the second set of users that generated the corresponding second set of labeling data. The medical scan labeling quality assurance system **3004** updates each of a second set of user profile entries in the user database for each corresponding one of the second set of users based on the second performance score data of the second set of performance score data assigned to the corresponding one of the second set of users. The medical scan labeling quality assurance system **3004** transmits each second performance score data of the second set

of performance score data to a corresponding one of the second set of client devices for display, via the first interactive interface, to a corresponding one of the second set of users to which the each second performance score data is assigned.

In various embodiments, the medical scan labeling quality assurance system **3004** identifies a subset of the set of users with performance score data that compares unfavorably to a performance score data upper bound. Selecting the second set of users includes selecting only ones of the set of users identified in the subset. In various embodiments, the medical scan labeling quality assurance system **3004** identifies a second subset of the plurality of users with performance score data that compares unfavorably to a performance score data upper bound. The second subset includes at least one user of the plurality of users that is not included in the set of users. Selecting the second set of users includes selecting only ones of the set of users identified in the second subset.

In various embodiments, the medical scan labeling quality assurance system **3004** identifies a second subset of the set of users with performance score data that compares favorably to the performance score data upper bound. The medical scan labeling quality assurance system **3004** updates user profile entries in the user database for each one of the second subset of the set of users to indicate an expert status for each one of the second subset of the set of users with in response to the performance score data that compares favorably to the performance score data upper bound. Selecting the expert user and the second expert user includes determining that the expert user and the second expert user have corresponding user profile entries indicating the expert status. In various embodiments, identifying the second subset of the set of users further includes determining for each of the set of users, a number of medical scans labeled with performance score data that compare favorably to the performance score data upper bound. Only ones of the set of users with a corresponding number that compares favorably to a threshold are included in the second set of users.

In various embodiments, the medical scan labeling quality assurance system **3004** identifies a subset of the set of users with performance score data that compare favorably to a performance score data lower bound. Selecting the second set of users includes selecting only ones of the set of users identified in the subset. The medical scan labeling quality assurance system **3004** identifies a second subset of the set of users with performance score data that compare unfavorably to a performance score data lower bound. The medical scan labeling quality assurance system **3004** transmits a notification to a subset of the set of client devices corresponding to each of the second subset of the set of users indicating a dismissal from labeling of future medical scans.

In various embodiments, the medical scan labeling quality assurance system **3004** selects the set of medical scans randomly. In various embodiments, the medical scan labeling quality assurance system **3004** selects the set of medical scans based on a set of determined selection requirements, where the set of determined selection requirements indicates at least one modality and further indicates at least one anatomical region. In various embodiments, the medical scan labeling quality assurance system **3004** selects the set of medical scans based on an urgency rating of the set of medical scans.

As may be used herein, the terms “substantially” and “approximately” provides an industry-accepted tolerance for its corresponding term and/or relativity between items. Such an industry-accepted tolerance ranges from less than one percent to fifty percent and corresponds to, but is not limited

to, component values, integrated circuit process variations, temperature variations, rise and fall times, and/or thermal noise. Such relativity between items ranges from a difference of a few percent to magnitude differences. As may also be used herein, the term(s) “configured to”, “operably coupled to”, “coupled to”, and/or “coupling” includes direct coupling between items and/or indirect coupling between items via an intervening item (e.g., an item includes, but is not limited to, a component, an element, a circuit, and/or a module) where, for an example of indirect coupling, the intervening item does not modify the information of a signal but may adjust its current level, voltage level, and/or power level. As may further be used herein, inferred coupling (i.e., where one element is coupled to another element by inference) includes direct and indirect coupling between two items in the same manner as “coupled to”. As may even further be used herein, the term “configured to”, “operable to”, “coupled to”, or “operably coupled to” indicates that an item includes one or more of power connections, input(s), output(s), etc., to perform, when activated, one or more its corresponding functions and may further include inferred coupling to one or more other items. As may further be used herein, the term “associated with”, includes direct and/or indirect coupling of separate items and/or one item being embedded within another item. As may still further be used herein, the term “automatically” refers to an action caused directly by a processor of a computer network in response to a triggering event and particularly without human interaction.

As may be used herein, the term “compares favorably”, indicates that a comparison between two or more items, signals, etc., provides a desired relationship. For example, when the desired relationship is that signal 1 has a greater magnitude than signal 2, a favorable comparison may be achieved when the magnitude of signal 1 is greater than that of signal 2 or when the magnitude of signal 2 is less than that of signal 1. As may be used herein, the term “compares unfavorably”, indicates that a comparison between two or more items, signals, etc., fails to provide the desired relationship.

As may also be used herein, the terms “processing module”, “processing circuit”, “processor”, “processing device” and/or “processing unit” may be a single processing device or a plurality of processing devices. Such a processing device may be a microprocessor, micro-controller, digital signal processor, graphics processing unit, microcomputer, central processing unit, field programmable gate array, programmable logic device, state machine, logic circuitry, analog circuitry, digital circuitry, and/or any device that manipulates signals (analog and/or digital) based on hard coding of the circuitry and/or operational instructions. The processing module, module, processing circuit, and/or processing unit may be, or further include, memory and/or an integrated memory element, which may be a single memory device, a plurality of memory devices, and/or embedded circuitry of another processing module, module, processing circuit, and/or processing unit. Such a memory device may be a read-only memory, random access memory, volatile memory, non-volatile memory, static memory, dynamic memory, flash memory, cache memory, and/or any device that stores digital information. Note that if the processing module, module, processing circuit, and/or processing unit includes more than one processing device, the processing devices may be centrally located (e.g., directly coupled together via a wired and/or wireless bus structure) or may be distributedly located (e.g., cloud computing via indirect coupling via a local area network and/or a wide area

network). Further note that if the processing module, module, processing circuit, and/or processing unit implements one or more of its functions via a state machine, analog circuitry, digital circuitry, and/or logic circuitry, the memory and/or memory element storing the corresponding operational instructions may be embedded within, or external to, the circuitry comprising the state machine, analog circuitry, digital circuitry, and/or logic circuitry. Still further note that, the memory element may store, and the processing module, module, processing circuit, and/or processing unit executes, hard coded and/or operational instructions corresponding to at least some of the steps and/or functions illustrated in one or more of the Figures and/or described herein. Such a memory device or memory element can be included in an article of manufacture. While the processing module, module, processing circuit, and/or processing unit device may be a general purpose computing device, the execution of the hard coded and/or operational instructions by the processing module, module, processing circuit, and/or processing unit configures such a general purpose computing device as a special purpose computing device to implement the corresponding steps and/or functions illustrated in one or more of the Figures and/or described herein. In particular, the hard coded and/or operational instructions by the processing module, module, processing circuit, and/or processing unit implement acts and algorithms performed by the processing module, module, processing circuit, and/or processing unit. Such acts and algorithms can be identified by name, can be illustrated via flowchart and/or described in words.

One or more embodiments have been described above with the aid of method steps illustrating the performance of specified functions and relationships thereof. The boundaries and sequence of these functional building blocks and method steps have been arbitrarily defined herein for convenience of description. Alternate boundaries and sequences can be defined so long as the specified functions and relationships are appropriately performed. Any such alternate boundaries or sequences are thus within the scope and spirit of the claims. Further, the boundaries of these functional building blocks have been arbitrarily defined for convenience of description. Alternate boundaries could be defined as long as the certain significant functions are appropriately performed. Similarly, flow diagram blocks may also have been arbitrarily defined herein to illustrate certain significant functionality.

To the extent used, the flow diagram block boundaries and sequence could have been defined otherwise and still perform the certain significant functionality. Such alternate definitions of both functional building blocks and flow diagram blocks and sequences are thus within the scope and spirit of the claims. One of average skill in the art will also recognize that the functional building blocks, and other illustrative blocks, modules and components herein, can be implemented as illustrated or by discrete components, application specific integrated circuits, processors executing appropriate software and the like or any combination thereof.

In addition, a flow diagram may include a “start” and/or “continue” indication. The “start” and “continue” indications reflect that the steps presented can optionally be incorporated in or otherwise used in conjunction with other routines. In this context, “start” indicates the beginning of the first step presented and may be preceded by other activities not specifically shown. Further, the “continue” indication reflects that the steps presented may be performed multiple times and/or may be succeeded by other activities not specifically shown. Further, while a flow diagram indi-

115

cates a particular ordering of steps, other orderings are likewise possible provided that the principles of causality are maintained.

The one or more embodiments are used herein to illustrate one or more aspects, one or more features, one or more concepts, and/or one or more examples. A physical embodiment of an apparatus, an article of manufacture, a machine, and/or of a process may include one or more of the aspects, features, concepts, examples, etc. described with reference to one or more of the embodiments discussed herein. Further, from figure to figure, the embodiments may incorporate the same or similarly named functions, steps, modules, etc. that may use the same or different reference numbers and, as such, the functions, steps, modules, etc. may be the same or similar functions, steps, modules, etc. or different ones.

The term "system" is used in the description of one or more of the embodiments. A system implements one or more functions via a device such as a processor or other processing device or other hardware that may include or operate in association with a memory that stores operational instructions. A system may operate independently and/or in conjunction with software and/or firmware. As also used herein, a system may contain one or more sub-system, each of which may be one or more systems.

As may further be used herein, a computer readable memory includes one or more memory elements. A memory element may be a separate memory device, multiple memory devices, or a set of memory locations within a memory device. Such a memory device may be a read-only memory, random access memory, volatile memory, non-volatile memory, static memory, dynamic memory, flash memory, cache memory, and/or any device that stores digital information. The memory device may be in a form a solid state memory, a hard drive memory, cloud memory, thumb drive, server memory, computing device memory, and/or other physical medium for storing digital information.

While particular combinations of various functions and features of the one or more embodiments have been expressly described herein, other combinations of these features and functions are likewise possible. The present disclosure is not limited by the particular examples disclosed herein and expressly incorporates these other combinations.

What is claimed is:

1. A medical scan system, comprising:

at least one processing system that includes a processor; and

at least one memory that stores executable instructions that, when executed by the processing system, cause the medical scan system to:

train a medical scan analysis function based on a training set, the medical scan function utilizing artificial intelligence;

transmit a set of medical scans to a set of client devices associated with a set of users;

receive a set of labeling data from each of the set of client devices, wherein each set of labeling data is generated by a corresponding one of the set of client devices, wherein each set of labeling data includes labeling data for each of the set of medical scans, and wherein the labeling data for each of the set of medical scans is generated by the corresponding one of the set of client devices in response to at least one prompt to provide the labeling data via a first interactive interface in conjunction with display of the each of the set of medical scans;

116

transmit the set of medical scans and each set of labeling data to an expert client device associated with an expert user that is not included in the set of users;

receive a set of golden labeling data and a plurality of sets of correction data from the expert client device, wherein the set of golden labeling data is generated by the expert client device and includes golden labeling data for each of the set of medical scans based on at least one prompt to provide the golden labeling data via a second interactive interface in conjunction with display of the set of medical scans, and wherein each set of correction data is generated by the expert client device in response to at least one additional prompt to provide the each of the set of correction data via the second interactive interface in conjunction with display of each corresponding set of labeling data; and

generate a set of performance score data by generating performance score data for each corresponding set of labeling data based on a corresponding set of correction data of the plurality of sets of correction data; transmit each correction data of the set of correction data to a corresponding one of the set of client devices for display;

assign each performance score data of the set of performance score data to a corresponding one of the set of users that generated the corresponding set of labeling data;

generate an updated training set that includes the set of golden labeling data;

retrain the medical scan analysis function utilizing the artificial intelligence based on the updated training set, wherein model parameters of the medical scan analysis function are updated, based on the updated training set, to improve performance of the medical scan analysis function; and

generate inference data for another medical scan by performing the medical scan analysis function utilizing the artificial intelligence on image data of the another medical scan, wherein the inference data indicates an inferred abnormality.

2. The medical scan system of claim 1, wherein the set of medical scans are transmitted to the set of client devices associated with the set of users based on elapsing of a scheduled interval.

3. The medical scan system of claim 1, wherein the set of users and the expert user are selected from a plurality of users of the medical scan system.

4. The medical scan system of claim 1, wherein each set of correction data includes comment data, and wherein the comment data is generated by the expert client device in response to text entered by the expert user in at least one text box displayed by the second interactive interface.

5. The medical scan system of claim 1, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to: transmit each correction data of the set of correction data to all of the set of client devices for display in conjunction with the set of medical scans.

6. The medical scan system of claim 1, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to: select a second set of users from a user database, wherein a set difference between the second set of users and the set of users is non-null; and

117

transmit the set of medical scans and the set of correction data to each one of a second set of client devices corresponding to the second set of users for display, via a third interactive interface, in conjunction with set of medical scans.

7. The medical scan system of claim 6, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to: determine a common error in the sets of labeling data for at least one of the set of medical scans based on the set of golden labeling data; and identify at least one set of labeling data that includes the common error; wherein only the at least one of the set of labeling data that includes the common error and only the at least one of the set of medical scans corresponding to the common error are transmitted to the second set of client devices, and wherein the third interactive interface indicates the common error.

8. The medical scan system of claim 7, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to: generate labeling commonality data by comparing the sets of labeling data to each other to determine similar ones of the sets of labeling data for at least one of the set of medical scans; and generate common error data based on comparing the labeling commonality data to the golden labeling data to determine a proper subset of the similar ones of the sets of labeling data that compares unfavorably to the golden labeling data; wherein the common error is determined based on the common error data.

9. The medical scan system of claim 7, wherein the common error is identified by the expert client device in response to at least one additional prompt to provide the common error data via the second interactive interface.

10. The medical scan system of claim 1, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to: determine that a scheduled interval has elapsed since transmission of the set of medical scans to the set of client devices; select a second set of users from a user database after the scheduled interval has elapsed since selection of the set of users in response to determining the scheduled interval has elapsed; select a second set of medical scans from a medical scan database in response to determining the scheduled interval has elapsed, wherein an intersection between the set of medical scans and the second set of medical scans is null;

transmit the second set of medical scans, via a network, to a second set of client devices associated with the second set of users, wherein the second set of medical scans are displayed to the second set of users via the first interactive interface displayed by a second set of display devices corresponding to the second set of client devices;

receive a second set of labeling data from each of the second set of client devices via the network, wherein each second set of labeling data is generated by a corresponding one of the second set of client devices, wherein each second set of labeling data includes second labeling data for each of the second set of medical scans, and wherein the second labeling data for each of the second set of medical scans is generated by

118

the corresponding one of the second set of client devices in response to at least one prompt to provide the second labeling data via the first interactive interface in conjunction with display of the each of the second set of medical scans;

select a second expert user from the user database, wherein the second expert user is not included in the second set of users;

transmit the second set of medical scans, via the network, to a second expert client device associated with the second expert user, wherein the second set of medical scans are displayed to the second expert user via a second interactive interface displayed by a second expert display device corresponding to the second expert user;

receive a second set of golden labeling data from the second expert client device via the network, wherein the second set of golden labeling data is generated by the second expert client device and includes second golden labeling data for each of the second set of medical scans, and wherein the second golden labeling data for each of the second set of medical scans is generated by the second expert client device in response to at least one prompt to provide the second golden labeling data via the second interactive interface in conjunction with display of the each of the second set of medical scans;

generate a second set of performance score data by generating second performance score data for each corresponding second set of labeling data by comparing second labeling data of each second set of labeling data to the second golden labeling data of the second set of global labeling data;

assign each second performance score data of the second set of performance score data to a corresponding one of the second set of users that generated the corresponding second set of labeling data;

update each of a second set of user profile entries in the user database for each corresponding one of the second set of users based on the second performance score data of the second set of performance score data assigned to the corresponding one of the second set of users; and transmit each second performance score data of the second set of performance score data to a corresponding one of the second set of client devices for display, via the first interactive interface, to a corresponding one of the second set of users to which the each second performance score data is assigned;

generate a second updated training set that includes the second set of golden labeling data; and

retrain the medical scan analysis function utilizing the artificial intelligence based on the second updated training set, wherein model parameters of the medical scan analysis function are updated, based on the second updated training set, to improve performance of the medical scan analysis function.

11. The medical scan system of claim 10, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to: identify a subset of the set of users with performance score data that compares unfavorably to a performance score data upper bound;

wherein selecting the second set of users includes selecting only ones of the set of users identified in the subset.

12. The medical scan system of claim 10, wherein the set of users are included in a plurality of users, and wherein the

119

executable instructions, when executed by the at least one processing system, further cause the medical scan system to:

identify a second subset of the plurality of users with performance score data that compares unfavorably to a performance score data upper bound, wherein the second subset includes at least one user of the plurality of users that is not included in the set of users;

wherein selecting the second set of users includes selecting only ones of the set of users identified in the second subset.

13. The medical scan system of claim **10**, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to:

identify a second subset of the set of users with performance score data that compares favorably to the performance score data upper bound; and

update user profile entries in the user database for each one of the second subset of the set of users to indicate an expert status for each one of the second subset of the set of users in response to the performance score data that compares favorably to the performance score data upper bound;

wherein selecting the expert user and the second expert user includes determining that the expert user and the second expert user have corresponding user profile entries indicating the expert status.

14. The medical scan system of claim **13**, wherein identifying the second subset of the set of users further includes:

determining for each of the set of users, a number of medical scans labeled with performance score data that compare favorably to the performance score data upper bound;

wherein only ones of the set of users with a corresponding number that compares favorably to a threshold are included in the second set of users.

15. The medical scan system of claim **10**, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to:

identify a subset of the set of users with performance score data that compare favorably to a performance score data lower bound, wherein selecting the second set of users includes selecting only ones of the set of users identified in the subset;

identify a second subset of the set of users with performance score data that compare unfavorably to a performance score data lower bound; and

transmit a notification to a subset of the set of client devices corresponding to each of the second subset of the set of users indicating a dismissal from labeling of future medical scans.

16. The medical scan system of claim **1**, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to select the set of medical scans randomly.

17. The medical scan system of claim **1**, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to select the set of medical scans based on a set of determined selection requirements, wherein the set of determined selection requirements indicates at least one modality and further indicates at least one anatomical region.

18. The medical scan system of claim **1**, wherein the executable instructions, when executed by the at least one processing system, further cause the medical scan system to select the set of medical scans based on an urgency rating of the set of medical scans.

120

19. A method comprising:

training a medical scan analysis function based on a training set, the medical scan function utilizing artificial intelligence;

transmitting a set of medical scans to a set of client devices associated with a set of users;

receiving a set of labeling data from each of the set of client devices, wherein each set of labeling data is generated by a corresponding one of the set of client devices, wherein each set of labeling data includes labeling data for each of the set of medical scans, and wherein the labeling data for each of the set of medical scans is generated by the corresponding one of the set of client devices in response to at least one prompt to provide the labeling data via a first interactive interface in conjunction with display of the each of the set of medical scans;

transmitting the set of medical scans and each set of labeling data to an expert client device associated with an expert user that is not included in the set of users;

receiving a set of golden labeling data and a plurality of sets of correction data from the expert client device, wherein the set of golden labeling data is generated by the expert client device and includes golden labeling data for each of the set of medical scans based on at least one prompt to provide the golden labeling data via a second interactive interface in conjunction with display of the set of medical scans, and wherein each set of correction data is generated by the expert client device in response to at least one additional prompt to provide the each of the set of correction data via the second interactive interface in conjunction with display of each corresponding set of labeling data; and

generating a set of performance score data by generating performance score data for each corresponding set of labeling data based on a corresponding set of correction data of the plurality of sets of correction data;

transmitting each correction data of the set of correction data to a corresponding one of the set of client devices for display;

assigning each performance score data of the set of performance score data to a corresponding one of the set of users that generated the corresponding set of labeling data;

generating an updated training set that includes the set of golden labeling data;

retraining the medical scan analysis function utilizing the artificial intelligence based on the updated training set, wherein model parameters of the medical scan analysis function are updated, based on the updated training set, to improve performance of the medical scan analysis function; and

generating inference data for another medical scan by performing the medical scan analysis function utilizing the artificial intelligence on image data of the another medical scan, wherein the inference data indicates an inferred abnormality.

20. A method comprising:

training a medical scan analysis function based on a training set, the medical scan function utilizing artificial intelligence;

transmitting a set of medical scans and a plurality of sets of labeling data to an expert client device associated with an expert user, wherein each set of labeling data was generated by a corresponding one of a set of client devices corresponding to a set of users, wherein each set of labeling data includes labeling data for each of the set of medical scans, and wherein the labeling data

121

for each of the set of medical scans is generated by the corresponding one of the set of client devices in response to at least one prompt to provide the labeling data in conjunction with display of the each of the set of medical scans; 5
receiving a set of golden labeling data and a plurality of sets of correction data from the expert client device, wherein the set of golden labeling data is generated by the expert client device and includes golden labeling data for each of the set of medical scans based on at least one prompt to provide the golden labeling data in conjunction with display of the set of medical scans, and wherein each set of correction data is generated by the expert client device in response to at least one additional prompt to provide the each of the set of correction data in conjunction with display of each corresponding set of labeling data; and 10
generating a set of performance score data by generating performance score data for each corresponding set of 15

122

labeling data based on a corresponding set of correction data of the plurality of sets of correction data;
assigning each performance score data of the set of performance score data to a corresponding one of the set of users that generated the corresponding set of labeling data;
generating an updated training set that includes the set of golden labeling data;
retraining the medical scan analysis function utilizing the artificial intelligence based on the updated training set, wherein model parameters of the medical scan analysis function are updated, based on the updated training set, to improve performance of the medical scan analysis function; and
generating inference data for another medical scan by performing the medical scan analysis function utilizing the artificial intelligence on image data of the another medical scan, wherein the inference data indicates an inferred abnormality.

* * * * *