

Claims

1. A method for detecting a plurality of visual findings in one or more anatomical images of a subject, comprising:
 - providing one or more anatomical images of the subject;
 - 5 inputting the one or more anatomical images into a first convolutional neural network (CNN) component of a primary neural network to output a feature vector;
 - computing an indication of a plurality of visual findings being present in at least one of the one or more anatomical images by a dense layer of the primary neural network that takes as input the feature vector and outputs an indication of whether each of the plurality of
 - 10 visual findings is present in at least one of the one or more anatomical images,
 - wherein the primary neural network is trained on a training dataset including, for each of a plurality of subjects, one or more anatomical images, and a plurality of labels associated with the one or more anatomical images and each of the respective visual findings, wherein the primary neural network is trained by evaluating performance of a
 - 15 plurality of neural networks in detecting the plurality of visual findings, wherein the performance evaluation comprises accounting for correlation between one or more pairs of the plurality of visual findings.
2. The method of claim 1 wherein the visual findings are radiological findings in anatomical images comprising one or more chest x-ray (CXR) images.
- 20 3. The method of claim 1 or 2 wherein accounting for correlation between one or more pairs of the plurality of visual findings comprises evaluating the performance of each of the plurality of neural networks using a testing dataset that comprises a subset of the training dataset, where the testing dataset is selected such that the correlation between the one or more pairs of the plurality of findings in the testing dataset satisfies one or more criteria
- 25 selected from:
 - the correlation between the one or more pairs of the plurality of findings in the validation dataset does not differ by more than a predetermined percentage from the corresponding correlation in the full training dataset; and
 - the correlation between the one or more pairs of the plurality of findings in the
 - 30 validation dataset does not exceed a predetermined threshold.
4. The method of any one of claims 1 to 3 wherein accounting for correlation between one or more pairs of the plurality of visual findings comprises evaluating the performance of each of the plurality of neural networks for each of the plurality of visual findings and at least

one negation pair class which comprises anatomical images where a first one of the plurality of visual findings is identified in the absence of a second one of the plurality of visual findings.

5. The method of any one of the preceding claims, wherein the primary neural network is further trained by evaluating the performance of a plurality of neural networks in detecting the plurality of visual findings and selecting one or more best performing neural networks.
6. The method of any one of the preceding claims wherein the plurality of visual findings is organised as a hierarchical ontology tree.
7. A method for detecting a plurality of visual findings in one or more anatomical images of a subject, comprising:
 - providing one or more anatomical images of a subject;
 - inputting the one or more anatomical images into a first convolutional neural network (CNN) component of a primary neural network to output a feature vector;
 - computing an indication of a plurality of visual findings being present in at least one of the one or more anatomical images by a dense layer of the primary neural network that takes as input the feature vector and outputs an indication of whether each of the plurality of visual findings is present in at least one of the one or more anatomical images,
 - wherein the primary neural network is trained on a training dataset including, for each of a plurality of subjects, one or more anatomical images, and a plurality of labels associated with the one or more anatomical images and each of the respective visual findings, wherein the plurality of visual findings is organised as a hierarchical ontology tree and the training comprises evaluating performance of the neural network at different levels of the hierarchy of the ontology tree.
8. The method of claim 7 wherein the visual findings are radiological findings in anatomical images comprising one or more chest x-ray (CXR) images.
9. The method of claim 7 or 8 wherein:
 - the hierarchical ontology tree comprises internal nodes and terminal leaves;
 - at least one of the plurality of labels is associated with a terminal leaf in the hierarchical ontology tree, and at least one of the plurality of labels is associated with an internal node in the hierarchical ontology tree; and
 - the primary neural network outputs an indication of whether each of the plurality of

visual findings is present in at least one of the one or more anatomical images, for visual findings that include internal nodes and terminal leaves.

10. The method of any one of claims 7 to 9 wherein the plurality of labels associated with at least a subset of the one or more anatomical images and each of the respective
5 visual findings in the training dataset may be derived from the results of review of the one or more anatomical images by at least one expert.

11. The method of claim 10 wherein review of the one or more anatomical images comprise by the at least one expert using a labelling tool that allows the expert to select labels presented in a hierarchical menu.

10 12. The method of any one of claims 7 to 11 wherein the primary neural network is trained by evaluating performance of a plurality of neural networks in detecting the plurality of visual findings, wherein the performance evaluation comprises accounting for correlation between one or more pairs of the plurality of visual findings.

13. The method of claim 12 wherein accounting for correlation between one or more
15 pairs of the plurality of visual findings comprises evaluating the performance of each of the plurality of neural networks for each of the plurality of visual findings and at least one negation pair class which comprises anatomical images where a first one of the plurality of visual findings is identified in the absence of a second one of the plurality of visual findings.

14. A method for detecting a plurality of visual findings in one or more anatomical
20 images of a subject, comprising:
providing one or more anatomical images of a subject;
inputting the one or more anatomical images into a first convolutional neural network (CNN) component of a primary neural network to output a feature vector;
computing an indication of a plurality of visual findings being present in at least one
25 of the one or more anatomical images by a dense layer of the primary neural network that takes as input the feature vector and outputs an indication of whether each of the plurality of visual findings is present in at least one of the one or more anatomical images,
wherein the primary neural network is trained on a training dataset including, for each of a plurality of subjects, one or more anatomical images, and a plurality of labels
30 associated with the one or more anatomical images and each of the respective visual findings, wherein the plurality of visual findings is organised as a hierarchical ontology tree, and

wherein the primary neural network is trained by evaluating the performance of a plurality of neural networks in detecting the plurality of visual findings and at least one negation pair class which comprises anatomical images where a first one of the plurality of visual findings is identified in the absence of a second one of the plurality of visual findings.

- 5 15. The method of claim 11 wherein the visual findings are radiological findings in anatomical images comprising one or more chest x-ray (CXR) images.
16. The method of claim 15 wherein the at least one negation pair class is selected from the group comprising: ('pneumothorax', 'subcutaneous_emphysema'), ('pneumothorax', 'intercostal_drain'), ('pneumothorax', 'tracheal_deviation'), ('pleural_effusion', 'intercostal_drain'), ('pleural_effusion', 'cardiomegaly'), ('significant_collapse', 'ett'), ('significant_collapse', 'diaphragmatic_elevation'), ('significant_collapse', 'tracheal_deviation'), ('significant_collapse', 'linear_atelectasis'), ('interstitial_thickening_volloss', 'linear_atelectasis'), ('interstitial_thickening_volloss_lower', 'linear_atelectasis'), ('interstitial_thickening_volloss_upper', 'interstitial_thickening_upper'), ('cavitating_mass', 'cavitating_mass_internal_content'), ('pneumomediastinum', 'subcutaneous_emphysema'), ('dish', 'spine_arthritis'), ('shoulder_dislocation', 'acute_humerus_fracture'), ('shoulder_dislocation', 'chronic_humerus_fracture'), ('rib_lesion', 'humeral_lesion'), ('rib_lesion', 'clavicle_lesion'), ('rib_lesion', 'scapular_lesion'), ('rib_lesion', 'spine_lesion'), ('clavicle_lesion', 'spine_lesion'), ('scapular_lesion', 'spine_lesion'), ('rib_resection', 'lung_sutures'), ('acute_humerus_fracture', 'chronic_humerus_fracture'), ('lung_lesion', 'surgical_clip'), ('lung_lesion', 'lung_sutures'), ('lung_lesion', 'lung_resection_volloss'), ('bullae', 'hyperinflation'), ('bullae', 'hyperlucency'), ('cardiomegaly', 'electronic_cardiac_devices'), ('cardiomegaly', 'cardiac_valve_prosthesis'), ('cardiomegaly', 'pulmonary_congestion'), ('cardiomegaly', 'sternotomy_wires'), ('cardiomegaly', 'airspace_opacity_without_focus'), ('cardiomegaly', 'interstitial_thickening_no_volloss'), ('acute_aortic_syndrome', 'tracheal_deviation'), ('aortic_arch_calcification', 'coronary_stent'), ('distended_bowel', 'subdiaphragmatic_gas'), ('airspace_opacity', 'pleural_effusion'), ('airspace_opacity', 'loculated_effusion'), ('ett', 'ngt'), ('ett', 'cvc'), ('ett', 'pac'), ('ett', 'intercostal_drain'), ('ngt', 'cvc'), ('ngt', 'pac'), ('ngt', 'intercostal_drain'), ('cvc', 'pac'), ('cvc', 'intercostal_drain'), ('pac', 'intercostal_drain'), ('kyphosis', 'scoliosis'), ('osteopaenia', 'spine_wedge_fracture'), and ('mastectomy', 'axillary_clips').
17. The method of any one of claims 14 to 16, wherein evaluating the performance of the primary neural network for each of the plurality of visual findings and at least one

negation pair class comprises computing a combined performance across the plurality of visual findings and the at least one negation pair class.

18. The method of any one of claims 14 to 17 wherein the primary neural network is further trained by evaluating performance of a plurality of neural networks in detecting the plurality of visual findings, wherein the performance evaluation comprises accounting for correlation between one or more pairs of the plurality of visual findings.

19. The method of any one of the preceding claims wherein the plurality of visual findings includes visual findings selected from Table 1 or Table 2.

20. The method of claim 19 wherein the plurality of visual findings includes at least 80 visual findings.

21. The method of any one of the preceding claims further comprising computing, by a decoder coupled to the CNN component of the primary neural network, a segmentation mask indicating a localisation for at least one of the plurality of visual findings.

22. The method of any one of the preceding claims wherein the one or more anatomical images comprise at least two of the anatomical images captured at different orientations of a body portion of the subject.

23. The method of claim 23 wherein the one or more anatomical images comprise at least one image that is captured by an imaging device when the subject is oriented anterior-posterior (AP) or posterior-anterior (PA) relative to the imaging device and at least one image that is captured by an imaging device when the subject is oriented laterally relative to the imaging device.

24. The method of claim 23 wherein the one or more anatomical images comprise at least one image that is captured by an imaging device when the subject is oriented anterior-posterior (AP) relative to the imaging device, at least one image that is captured by an imaging device when the subject is oriented poster-anterior (PA) relative to the imaging device, and at least one image that is captured by an imaging device when the subject is oriented laterally relative to the imaging device.

25. The method of claims 23 or 24 further comprising:
inputting an anatomical image of the at least two anatomical images captured at different orientations into a further CNN component of a further neural network to output a further feature vector; and

inputting the further feature vector into a dense layer of the further neural network to generate an indication of orientation of the input anatomical image.

26. The method of claim 25 wherein the primary neural network further comprises at least a second CNN component, and the method further comprises:

5 inputting a first anatomical image of the at least two anatomical images captured at different orientations into the first CNN component of the primary neural network to output a first feature vector;

inputting a second anatomical image of the at least two anatomical images captured at different orientations into the second CNN component of the primary neural network to
10 output a second feature vector; and

inputting a feature vector that combines the first feature vector and the second feature vector into the dense layer of the primary neural network.

27. The method of claim 26 wherein the at least two anatomical images captured at different orientations comprise a third anatomical image, the primary neural network further
15 comprises a third CNN component, and the method further comprises:

inputting the third anatomical image into the third CNN component of the primary neural network to output a third feature vector; and

inputting a feature vector that combines the first feature vector, the second feature vector and the third feature vector into the dense layer of the primary neural network.

20 28. The method of claim 26 wherein the first and second CNN components comprise a shared CNN component.

29. The method of claim 27 wherein at least two of the first, second, and third CNN components comprise a shared CNN component.

30. A method for training a neural network to detect a plurality of visual findings in one
25 or more anatomical images of a subject, comprising:

providing a primary neural network comprising a first convolutional neural network (CNN) component that takes as input one or more anatomical images of a subject and outputs a first feature vector, and a dense layer that takes as input a feature vector comprising the first feature vector and outputs an indication of whether each of the plurality
30 of visual findings is present in at least one of the one or more anatomical images;

retrieving, from a data store, a training dataset including, for each of a plurality of subjects, one or more anatomical images, and a plurality of labels associated with the one or

more anatomical images and each of the respective visual findings; and
 training the primary neural network using the training dataset,
 wherein the training comprises evaluating performance of the primary neural
 network in detecting the plurality of visual findings relative to one or more similar neural
 5 networks, wherein the performance evaluation comprises accounting for correlation between
 one or more pairs of the plurality of visual findings.

31. The method of claim 30 wherein accounting for correlation between one or more
 pairs of the plurality of visual findings comprises evaluating the performance of each of the
 plurality of neural networks using a testing dataset that comprises a subset of the training
 10 dataset, where the testing dataset is selected such that the correlation between the one or
 more pairs of the plurality of findings in the testing dataset satisfies one or more criteria
 selected from:

the correlation between the one or more pairs of the plurality of findings in the
 validation dataset does not differ by more than a predetermined percentage from the
 15 corresponding correlation in the full training dataset; and
 the correlation between the one or more pairs of the plurality of findings in the
 validation dataset does not exceed a predetermined threshold.

32. The method of claim 30 or 31 wherein accounting for correlation between one or
 more pairs of the plurality of visual findings comprises evaluating the performance of each of
 20 the plurality of neural networks for each of the plurality of visual findings and at least one
 negation pair class which comprises anatomical images where a first one of the plurality of
 visual findings is identified in the absence of a second one of the plurality of visual findings.

33. The method of any one of claims 30 to 32, wherein the training comprises
 evaluating the performance of a plurality of neural networks in detecting the plurality of visual
 25 findings and selecting one or more best performing neural networks.

34. The method of any one of claims 30 to 33 wherein the plurality of visual findings is
 organised as a hierarchical ontology tree.

35. A method for training a neural network to detect a plurality of visual findings in one
 or more anatomical images of a subject, comprising:
 30 providing a primary neural network comprising a first convolutional neural network
 (CNN) component that takes as input one or more anatomical images of a subject and
 outputs a first feature vector, and a dense layer that takes as input a feature vector

comprising the first feature vector and outputs an indication of whether each of the plurality of visual findings is present in at least one of the one or more anatomical images;

retrieving, from a data store, a training dataset including, for each of a plurality of subjects, one or more anatomical images, and a plurality of labels associated with the one or
5 more anatomical images and each of the respective visual findings; and

training the primary neural network using the training dataset,
wherein the plurality of visual findings is organised as a hierarchical ontology tree,
and the training comprises evaluating performance of the primary neural network at different levels of the hierarchy of the ontology tree.

10 36. The method of claim 35 wherein:

the hierarchical ontology tree comprises internal nodes and terminal leaves;
at least one of the plurality of labels is associated with a terminal leaf in the hierarchical ontology tree, and at least one of the plurality of labels is associated with an internal node in the hierarchical ontology tree; and

15 the primary neural network outputs an indication of whether each of the plurality of visual findings is present in at least one of the one or more anatomical images, for visual findings that include internal nodes and terminal leaves.

37. The method of claim 35 or 36 wherein the plurality of labels associated with at least a subset of the one or more anatomical images and each of the respective visual findings in
20 the training dataset may be derived from the results of review of the one or more anatomical images by at least one expert.

38. The method of claim 37 wherein review of the one or more anatomical images comprises, by the at least one expert, using a labelling tool that allows the expert to select labels presented in a hierarchical menu.

25 39. The method of any one of claims 35 to 37 wherein the training comprises evaluating performance of a plurality of neural networks in detecting the plurality of visual findings, wherein the performance evaluation comprises accounting for correlation between one or more pairs of the plurality of visual findings.

40. The method of claim 39 wherein accounting for correlation between one or more
30 pairs of the plurality of visual findings comprises evaluating the performance of each of the plurality of neural networks for each of the plurality of visual findings and at least one

negation pair class which comprises anatomical images where a first one of the plurality of visual findings is identified in the absence of a second one of the plurality of visual findings.

41. A method for training a neural network to detect a plurality of visual findings in one or more anatomical images of a subject, comprising:

- 5 providing a primary neural network comprising a first convolutional neural network (CNN) component that takes as input one or more anatomical images of a subject and outputs a first feature vector, and a dense layer that takes as input a feature vector comprising the first feature vector and outputs an indication of whether each of the plurality of visual findings is present in at least one of the one or more anatomical images;
- 10 retrieving, from a data store, a training dataset including, for each of a plurality of subjects, one or more anatomical images, and a plurality of labels associated with the one or more anatomical images and each of the respective visual findings; and
- training the primary neural network using the training dataset, wherein the plurality of visual findings is organised as a hierarchical ontology tree,
- 15 and the training comprises evaluating performance of the primary neural network at different levels of the hierarchy of the ontology tree, and
- wherein the training further comprises evaluating performance of the primary neural network in detecting the plurality of visual findings, and at least one negation pair class which comprises anatomical images where a first one of the plurality of visual findings is
- 20 identified in the absence of a second one of the plurality of visual findings, relative to one or more similar neural networks.

42. The method of claim 41 wherein evaluating the performance of the primary neural network for each of the plurality of visual findings and at least one negation pair class comprises computing a combined performance across the plurality of visual findings and the

25 at least one negation pair class.

43. The method of claim 41 or 42 wherein the primary neural network is further trained by evaluating performance of a plurality of neural networks in detecting the plurality of visual findings, wherein the performance evaluation comprises accounting for correlation between one or more pairs of the plurality of visual findings.

- 30 44. A system for detecting a plurality of visual findings in one or more anatomical images of a subject, comprising:
 - at least one processor; and
 - at least one computer readable storage medium, accessible by the processor,

comprising instructions that, when executed by the processor, cause the processor to execute a method according to any one of claims 1 to 29.

45. A system for training a neural network to detect a plurality of visual findings in one or more anatomical images of a subject, comprising:

- 5 at least one processor; and
 at least one computer readable storage medium, accessible by the processor,
comprising instructions that, when executed by the processor, cause the processor to execute a method according to any one of claims 30 to 43.

46. A non-transitory computer readable storage media comprising instructions that,
10 when executed by at least one processor, cause the processor to execute a method according to any one of claims 1 to 43.