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;

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 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.

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.

3. The method of claim 1 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 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 validation dataset does not exceed a predetermined threshold.

4. 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.

5. The method of claim 4 wherein the visual findings are radiological findings in anatomical images comprising one or more chest x-ray (CXR) images.

6. The method of claim 4 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.

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

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.

8. The method of claim 1 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.

9. The method of claim 8 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.

10. The method of claim 8 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.

11. The method of claim 10 wherein the primary neural network further comprises at least a second CNN component, and the method further comprises:

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 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.

12. The method of claim 11 wherein the at least two anatomical images captured at different orientations comprise a third anatomical image, the primary neural network further 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.

13. The method of claim 11 wherein the first and second CNN components comprise a shared CNN component.

14. The method of claim 12 wherein at least two of the first, second, and third CNN components comprise a shared CNN component.

15. A method for training a neural network to detect a plurality of visual findings in one 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 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 networks, wherein the performance evaluation comprises accounting for correlation between one or more pairs of the plurality of visual findings.

16. The method of claim 15 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 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 validation dataset does not exceed a predetermined threshold.

17. The method of claim 15 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.

18. A method for training a neural network to detect a plurality of visual findings in one 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 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 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.

19. The method of claim 18 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.

20. The method of claim 18 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 the training dataset may be derived from the results of review of the one or more anatomical images by at least one expert.

21. The method of claim 20 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.

22. A method for training a neural network to detect a plurality of visual findings in one 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 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 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, 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 identified in the absence of a second one of the plurality of visual findings, relative to one or more similar neural networks.

23. The method of claim 22 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.

24. The method of claim 22 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.