

CLAIMS

What is claimed is:

1. A method comprising:

obtaining a set of classification models, each of which is configured to predict a label for an image, wherein the predicted label indicates a class of the image, 5
wherein each classification model of the set of classification models is configured to predict a label from a same set of labels;

applying the set of classification models on a calibration dataset of images thereby providing for each image of the calibration dataset of images an array of predicted labels and providing a set of arrays of predicted labels for the calibration 10
dataset;

calculating a disagreement measurement of the set of classification models on the calibration dataset, wherein the disagreement measurement is calculated based on the set of arrays of predicted labels for the calibration dataset, wherein the disagreement measurement is affected by differences between predictions of 15
classification models of the set of classification models;

applying the set of classification models on a production dataset of images thereby providing a set of arrays of predicted labels for the production dataset;

calculating a production disagreement measurement of the set of classification models on the production dataset; 20

determining a similarity measurement between the production disagreement measurement and the disagreement measurement; and

in response to the similarity measurement being below a predetermined threshold, indicating a data drift in the production dataset.

2. The method of claim 1, further comprising applying a predictor on the production 25
dataset, wherein the predictor is configured to predict, for an image, a label from the same set of labels.

3. The method of claim 2, wherein the set of classification models excludes the predictor.

4. The method of claim 2, wherein the set of classification models comprises the predictor.
5. The method of claim 2, further comprising, in response to said indicating the drift in the production dataset, retraining the predictor based on at least a portion of the production dataset. 5
6. The method of claim 5, further comprising re-computing the disagreement measurement based on the at least a portion of the production dataset.
7. The method of claim 1, wherein the disagreement measurement is computed based on a portion of images for which a tuple of classification models of the set of classification models provide different labels. 10
8. The method of claim 1, wherein the disagreement measurement is computed based on a portion of images for which a tuple of classification models of the set of classification models provide a same label.
9. The method of claim 1, wherein the set of classification models comprise n classification models denoted as $cls_1, \dots, cls_n$, wherein the disagreement measurement is computed based on a portion of images for which the set of classification models provide an array of predicted labels having a value of $(l_1, \dots, l_n)$, where l_i is a label from the same set of labels that is predicted by cls_i , wherein the value of $(l_1, \dots, l_n)$ is a heterogeneous value. 15
10. A computer program product comprising a non-transitory computer readable storage medium retaining program instructions, the computer program product configured to carry out the method of claim 1. 20
11. A method comprising:
 - obtaining an image dataset, wherein the image dataset comprises an ordered plurality of sets of images, wherein each set of images comprises images obtained over a time interval, wherein the sets of images are ordered according to their respective time intervals; 25

determining a training dataset, wherein the training dataset comprises images from the image dataset, wherein said determining the training set comprises:

determining weights for the ordered plurality of sets of images, wherein each set of images is associated with a weight, wherein at least two of the sets of images are associated with different weights; and

5

selecting from the ordered plurality of sets of images a subset of images to be comprised by the training dataset, wherein said selecting is based on the weights; and

training a classification model using the training dataset.

12. The method of claim 11, wherein a last set of images according to the ordered plurality of sets of images is a production dataset, wherein the production dataset is obtained during application of the classification model in production on unlabelled images.

10

13. The method of claim 11, wherein the weights are percentages, whereby each set of images is associated with a percentage, wherein said selecting the subset of images is performed by selecting from each set of images a portion of the set of images according to a ratio defined by the percentage associated with the each set.

15

14. The method of claim 11, wherein said determining percentages of images of each selected set of images to be comprised by the training dataset, comprises:

selecting percentages correlated with order of the time intervals, whereby lower percentages are determined for older training datasets.

20

15. The method of claim 11, wherein said determining the weights for the ordered plurality of sets of images comprises determining the weights according to a monotonic function.

16. The method of claim 15, wherein the monotonic function is a non-uniform function selected from:

25

a strictly increasing monotonic function;

a strictly decreasing monotonic function;

a linear function; or

a non-linear function.

17. The method of claim 11, wherein said determining the weights for the ordered plurality of sets of images comprises:

training a first classifier for a first set of images, wherein the first classifier is trained based on a first dataset, wherein the first dataset comprises the first set of images; 5

training a second classifier for a second set of images, wherein the second classifier is trained based on a second dataset, wherein the second dataset consists of the first dataset and the second set of images; wherein the second set of images is ordered immediately after the first set of images according to the ordered plurality of sets of images; and 10

determining the weight of the second set of images based on a difference in an accuracy measurement of the second classifier compared to an accuracy measurement of the first classifier.

18. The method of claim 17, wherein the difference in the accuracy measurement of the second classifier compared to the accuracy measurement of the first classifier indicates an improvement, wherein the weight of the second set of images is correlated to a measurement of the improvement. 15

19. The method of claim 17, wherein the difference in the accuracy measurement of the second classifier compared to the accuracy measurement of the first classifier indicates a reduction in accuracy, wherein the weight of the second set of images is a lowest weight value, wherein said selecting the subset of images comprises avoiding selecting any image from a set of images that is assigned the lowest weight value thereby dropping the second set of images. 20

20. The method of claim 11, wherein said determining the weights for the ordered plurality of sets of images is performed based on a function $f(x)$, wherein the method further comprises: 25

determining validation weights for the ordered plurality of sets of images, wherein the validation weights are determined based on a function $g(x)$, wherein $g(x) = -f(x)$;

selecting from the ordered plurality of sets of images, a validation dataset of images, wherein said selecting the validation dataset of images is based on the validation weights; and

utilizing the validation dataset of images to validate the classification model.

21. The method of claim 11, further comprises:

determining a first sub-series of sets of images from the ordered plurality of sets of images, wherein the first sub-series comprises a predetermined number of successive sets of images;

training, using the first sub-series, a classifier that is configured to predict a label for an image, wherein the label is indicative of a class of the image;

applying the classifier on a second sub-series of sets of images from the ordered plurality of sets of images, wherein the second sub-series has a number of sets, wherein a first set in the second sub-series is ordered, in accordance with the ordered plurality of sets of images, immediately after a last set in the first sub-series;

determining an accuracy measurement of the classifier on the second sub-series; and

wherein said determining weights for the ordered plurality of sets of images is based on the accuracy measurement.

22. The method of claim 21, wherein said determining weights for the ordered plurality of sets of images comprises assigning a lowest weight value for the sets of the second sub-series based on the accuracy measurement being above a threshold, wherein said selecting the subset of images comprises avoiding selecting any image from a set of images that is assigned the lowest weight value thereby dropping the second sub-series.

23. The method of claim 21, wherein said determining weights for the ordered plurality of sets of images comprises assigning a lowest weight value for a plurality of sequences of sets of images, wherein each sequence of the plurality of sequences has the number of sets.
24. A computer program product comprising a non-transitory computer readable storage medium retaining program instructions, the computer program product configured to carry out the method of claim 11. 5
25. A computerized apparatus having a processor, the processor being adapted to perform the steps of:
- obtaining a set of classification models, each of which is configured to predict a label for an image, wherein the predicted label indicates a class of the image, wherein each classification model of the set of classification models is configured to predict a label from a same set of labels; 10
 - applying the set of classification models on a calibration dataset of images thereby providing for each image of the calibration dataset of images an array of predicted labels and providing a set of arrays of predicted labels for the calibration dataset; 15
 - calculating a disagreement measurement of the set of classification models on the calibration dataset, wherein the disagreement measurement is calculated based on the set of arrays of predicted labels for the calibration dataset, wherein the disagreement measurement is affected by differences between predictions of classification models of the set of classification models; 20
 - applying the set of classification models on a production dataset of images thereby providing a set of arrays of predicted labels for the production dataset;
 - calculating a production disagreement measurement of the set of classification models on the production dataset; 25
 - determining a similarity measurement between the production disagreement measurement and the disagreement measurement; and

in response to the similarity measurement being below a predetermined threshold, indicating a data drift in the production dataset.

26. A computerized apparatus having a processor, the processor being adapted to perform the steps of:

obtaining an image dataset, wherein the image dataset comprises an ordered plurality of sets of images, wherein each set of images comprises images obtained over a time interval, wherein the sets of images are ordered according to their respective time intervals; 5

determining a training dataset, wherein the training dataset comprises images from the image dataset, wherein said determining the training set comprises: 10

determining weights for the ordered plurality of sets of images, wherein each set of images is associated with a weight, wherein at least two sets of images are associated with different weights; and

selecting from the ordered plurality of sets of images a subset of images to be comprised by the training dataset, wherein said selecting is based on the weights; and 15

training a classification model using the training dataset.