



(51) International Patent Classification:

G16H 50/20 (2018.01)

(21) International Application Number:

PCT/GB2020/050469

(22) International Filing Date:

27 February 2020 (27.02.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

1904481.7

29 March 2019 (29.03.2019)

GB

(71) Applicant: **ARM LIMITED** [GB/GB]; 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB).

(72) Inventors: **ÖZER, Emre**; c/o Arm Limited, 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB).

BROWN, Gavin; c/o Arm Limited, 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB). **REYNOLDS, Charles Edward Michael**; c/o Arm Limited, 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB). **KUFEL, Jędrzej**; c/o Arm Limited, 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB).

(74) Agent: **MALDEN, Nicholas**; D Young & Co LLP, 120 Holborn, London EC1N 2DY (GB).

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

(54) Title: FEATURE DATASET CLASSIFICATION




FIG. 7

(57) Abstract: Apparatuses and methods of operating such apparatuses are disclosed. An apparatus comprises feature dataset input circuitry to receive a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits. Class retrieval circuitry is responsive to reception of the feature dataset from the feature dataset input circuitry to retrieve from class indications storage a class indication for each feature data value received in the feature dataset, wherein class indications are predetermined and stored in the class indications storage for each permutation of the set of bits for each feature. Classification output circuitry is responsive to reception of class indications from the class retrieval circuitry to determine a classification in dependence on the class indications. A predicated class may thus be accurately generated from a simple apparatus.



SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR,
TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

Published:

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

FEATURE DATASET CLASSIFICATION

The present techniques relate to the field of data processing.

5 Various approaches may be taken to classify an input data set on the basis of a number of feature data values which make up that feature data set. For example an apparatus may be constructed on the basis of naïve Bayes classifiers which apply Bayes' Theorem. A common implementation is based on the Gaussian naïve Bayes algorithm, in which each factor of the likelihood term in the Bayes equation is
10 modelled as a (univariate) Gaussian distribution. A naïve Bayes algorithm implementation can be trained using training data sets (in which the desired class to be predicted is known) and then this trained model can be used for new input data sets to generate class predictions. Such an implementation may, in hardware, still require a significant level of computation ability in order to process each input data set and
15 generate the predicted class on the basis of the trained model. There may be some implementation contexts in which it is desirable for the class prediction to be generated by a simpler device with limited data processing ability.

 At least some examples provide an apparatus comprising: feature dataset input
20 circuitry to receive a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits; class retrieval circuitry responsive to reception of the feature dataset from the feature dataset input circuitry to retrieve from class indications storage a class indication for each feature data value received in the feature dataset, wherein class indications are
25 predetermined and stored in the class indications storage for each permutation of the set of bits for each feature; and classification output circuitry responsive to reception of class indications from the class retrieval circuitry to determine a classification in dependence on the class indications.

30 At least some examples provide a method of operating an apparatus comprising: receiving at a feature dataset input a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is

represented by a set of bits; retrieving from class indications storage a class indication for each feature data value received in the feature dataset, wherein class indications are predetermined and stored in the class indications storage for each permutation of the set of bits for each feature; and determining a classification in dependence on the class
5 indications.

At least some examples provide an apparatus comprising: means for receiving a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits; means for retrieving
10 from means for storing class indications a class indication for each feature data value received in the feature dataset, wherein class indications are predetermined and stored in the means for storing class indications for each permutation of the set of bits for each feature; and means for determining a classification in dependence on the class indications.

15

The present techniques will be described further, by way of example only, with reference to embodiments thereof as illustrated in the accompanying drawings, to be read in conjunction with the following description, in which:

20 Figure 1 schematically illustrates an apparatus according to some example embodiments;

Figure 2 schematically illustrates a training arrangement according to which a Gaussian naïve Bayes model implementation is trained using feature training data sets in some example embodiments;

25 Figures 3A and 3B show examples feature distributions on which the trained model may be based in some example embodiments;

Figure 4 schematically illustrates an apparatus comprising multiple class lookup tables which are accessed in parallel in some example embodiments;

Figure 5 schematically illustrates an apparatus comprising a single class lookup
30 table which is serially accessed for multiple feature values in some example embodiments;

Figure 6A schematically illustrates the weighting of class values in some example embodiments;

Figure 6B schematically illustrates a selection between indicated classes by vote in some example embodiments;

5 Figure 7 schematically illustrates a low precision implementation for receiving 5-bit feature values to be used to lookup 3-bit class values in some example embodiments;

Figure 8A schematically illustrates a set of sensors generating the input feature data set for an apparatus in some example embodiments;

10 Figure 8B schematically illustrates an apparatus embodied as a plastic fabricated device in some example embodiments;

Figure 9 schematically illustrates an apparatus designed to be a wearable device in some example embodiments; and

15 Figure 10 shows a sequence of steps which are taken according to the method of some example embodiments.

In one example herein there is an apparatus comprising: feature dataset input circuitry to receive a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits; class
20 retrieval circuitry responsive to reception of the feature dataset from the feature dataset input circuitry to retrieve from class indications storage a class indication for each feature data value received in the feature dataset, wherein class indications are predetermined and stored in the class indications storage for each permutation of the set of bits for each feature; and classification output circuitry responsive to reception
25 of class indications from the class retrieval circuitry to determine a classification in dependence on the class indications.

Before discussing features of the present techniques, for context the core features of the Naïve Bayes algorithm are first outlined. Naïve Bayes is a probabilistic
30 machine learning algorithm based on the application of Bayes' Theorem. It uses the simplifying assumption that all features are statistically independent once conditioned on the value of the class label. This simplifies the Bayes equation into the following:

$$y^* = \arg \max_{y \in Y} \frac{p(y) \prod_{i=1}^d p(x_i | y)}{p(x)}$$

Equation 1

5 y^* is the class label that maximises the equation where $Y = \{\text{set of classes}\}$, and d is the number of features, and x_i are the observed feature values. $p(y)$ are the priors, and $\prod_{i=1}^d p(x_i | y)$ is the likelihood function. The denominator term may be omitted, because it has no effect on the maximum, and the logarithm is taken to turn multiplications into additions:

10

$$y^* = \arg \max_{y \in Y} \left[\log p(y) + \sum_{i=1}^d \log p(x_i | y) \right].$$

Equation 2

Gaussian Naïve Bayes is a variant in which each factor of the likelihood term is
15 modelled as a (univariate) Gaussian distribution:

$$p(x|y) = \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

Equation 3

20 If the features are assumed to follow the Gaussian distribution, then each likelihood term will be substituted by the Gaussian probability density function as follows:

$$y^* = \arg \max_{y \in Y} \left[\log p(y) + \sum_{i=1}^d \left[-\log \sigma_{y,i} \sqrt{2\pi} - \frac{x_i - \mu_{y,i}}{2\sigma_{y,i}^2} \right] \right].$$

Equation 4

Eventually, this reduces to a simple form like the equation below:

$$y^* = \arg \max_{y \in Y} [C_y + \sum_{i=1}^d [K0_{y,i} + K1_{y,i} \cdot x_i + K2_{y,i} \cdot x_i^2]]$$

5

Equation 5

where C_y is the log prior of the class, and $K0_{y,i}$, $K1_{y,i}$ and $K2_{y,i}$ are constant and coefficients for each class/feature combination. Generally, although these values can be pre-computed in a training stage, and stored in memory, calculating the class probabilities is still compute-intensive requiring multiple MAC operations.

In this context the present techniques provide an apparatus which receives a feature data set comprising multiple feature data values and on the basis of that received feature data set determines a classification (i.e. a class) representative of that feature data set. However, instead of accumulating all of the features to find the probability of each possible class, and then determining the class with the maximum probability, an approach is proposed in which class probabilities for each possible value of each feature are precomputed. That is to say in the classifier apparatus proposed, instead of considering a single classifier with multiple features, multiple distinct classifiers are generated and from those one is selected to be the representative class. The inventors of the present techniques have established that in an apparatus implementing this approach this can enable the gate count to be reduced, and potentially make the operation of class determination faster.

In some embodiments the class indications stored in the class indications storage for each feature are each predetermined as a best class indication which maximises a Bayes Classifier for the feature in a training phase using feature training datasets. Accordingly, the class probabilities for each feature are precomputed in the training phase and the best class for that feature can then be selected under its corresponding Bayes classifier:

$$\hat{y}_i = \arg \max_{y \in Y} [C_y + K0_{y,i} + K1_{y,i} \cdot x_i + K2_{y,i} \cdot x_i^2],$$

Equation 6

Various forms of Bayes classifier may be implemented, but in some
 5 embodiments the Bayes Classifier is a Gaussian naïve Bayes Classifier.

Each feature of the set of features may be modelled using a range of different
 distribution types. In some embodiments the Bayes Classifier is based on a single
 distribution type used for each feature of the set of features. In some embodiments the
 10 Bayes Classifier is based on heterogeneous distribution types used for the set of
 features. These distribution types may take a variety of forms such as for example
 Gaussian, exponential, uniform and so on. The approach proposed does not constrain
 each feature to come from a particular type of distribution or for all features to come
 from the same type of distribution, which allows more flexibility in the
 15 implementation, for example if exact distributions of the features are known.

The retrieval of the class indications from the class indication storage may take
 a variety of forms but in some embodiments the class indications storage has a look-up
 table format and the class retrieval circuitry is arranged to perform a look-up procedure
 20 with respect to the look-up table format for each feature data value received in the
 feature dataset. This allows for a ready retrieval of precomputed class indications.

In some embodiments the class retrieval circuitry is arranged to retrieve in
 parallel the class indications for each feature data value received in the feature dataset.
 25 In some embodiments the class retrieval circuitry is arranged to retrieve in a serial
 sequence the class indications for each feature data value received in the feature
 dataset. Accordingly, it may be selected between the two different approaches, in
 dependence on the relative priority in a given implementation of the greater storage
 required when class indications are to be retrieved in parallel versus the longer
 30 retrieval time required for the class indications to be retrieved in a serial sequence.

Once the class indications have been retrieved from the class indications storage in the class retrieval circuitry, the final classification may be determined in a variety of ways, but in some embodiments the classification output circuitry is responsive to reception of the class indications from the class retrieval circuitry to
5 determine the classification by a vote amongst the class indications. The vote itself may have a variety of configurations, but for example the class selected may be that which is the most frequent class amongst the class indications retrieved by the class retrieval circuitry. In other words it may be the statistical mode of the set of class indications.

10

In some embodiments the class indications are weighted. This allows for a further degree of control over the final classification selected. This weighting may be predetermined in the sense that it is pre-calculated, for example where weights for the class indications are determined in a training phase using feature training datasets.
15 Alternatively these weights may be independently user defined. This allows greater user control over the allocation and selection of the classes.

In some embodiments weights of the class indications are used as a tiebreaker when the vote selects more than one class indication. Thus where selection between
20 the class indications to determine a unique class indication is not possible on the basis of the vote, for example because more than one class has been selected the same number of times in the process, then weights of the class indications may be used as tiebreaker in this situation to decide between them.

25 The inventors of the present techniques have found that successful implementations, in the sense that they maintain a usefully high prediction accuracy for the predicted class of a given input future data sets, can be maintained even when each feature data value is represented at a low precision. For example in some embodiments each feature data value is represented by a set of fewer than 10 bits.
30 Further, in some embodiments each feature data value is represented by a set of 5 bits.

In some embodiments the class indications are stored in the class indications storage using a representation of fewer than 5 bits. In some embodiments the class indications are stored in the class indication storage using a representation of 3 bits (i.e. allowing 8 different classes to be defined).

5

The present techniques may find applicability in a wide variety of contexts, but where they may be implemented in a notably low-complexity manner (in particular in terms of the gate count required) the techniques may find implementation in portable and indeed in wearable contexts. Accordingly, in some embodiments the apparatus is a wearable device.

10

The feature data set input may be provided from a variety of sources, but in some embodiments the feature dataset input is coupled to a plurality of sensors each providing a respective feature data value.

15

In some embodiments the apparatus is embodied as a plastic fabricated device. Such a data processing device, embodied in plastic (as opposed to for example being embodied as a silicon-based device) may make it particularly suitable for implementation as a wearable device, whether embedded in clothing or worn next to the skin for example.

20

The above mentioned low gate count of the apparatus may for example be noted in some embodiments in which the apparatus is embodied comprising fewer than 1000 logic gates. Further, in some embodiments the apparatus is embodied comprising fewer than 500 logic gates.

25

In one example herein there is a method of operating an apparatus comprising: receiving at a feature dataset input a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits; retrieving from class indications storage a class indication for each feature data value received in the feature dataset, wherein class indications are predetermined

30

and stored in the class indications storage for each permutation of the set of bits for each feature; and determining a classification in dependence on the class indications.

In one example herein there is an apparatus comprising: means for receiving a
5 feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits; means for retrieving from means for storing class indications a class indication for each feature data value received in the feature dataset, wherein class indications are predetermined and stored in the means for storing class indications for each permutation of the set of bits for
10 each feature; and means for determining a classification in dependence on the class indications.

Some particular embodiments are now described with reference to the figures.

15 Figure 1 schematically illustrates an apparatus 100 in some example embodiments. The apparatus comprises feature dataset input circuitry 101 which receives a set of feature data values. Four feature data values are shown in Figure 1 (and for simplicity and clarity of illustration this example number of inputs is continued through various example embodiments illustrated and discussed here), but
20 the present techniques are not limited to this number of inputs in the feature dataset. This feature data set comprising multiple feature data values is passed to class retrieval circuitry 102 which uses the individual feature data values to retrieve a corresponding set of class indications from class indications storage 103. These class indications are then passed to classification output circuitry 104 which determines a final, single
25 predicated classification based on that set of class indications received from the class retrieval circuitry 102. The final classification may be output from the apparatus 100 (as shown in Figure 1), but in other embodiments this classification may be utilised within the apparatus 100, for example to generate an indication which a user can perceive. The class indications stored in the class indications storage are
30 predetermined during a training phase for the apparatus, during which feature training data sets are used. This will be described in more detail with respect to Figure 2.

Figure 2 schematically illustrates a process for carrying out a training phase according to which a model (specifically the machine learning algorithm of the present techniques based on a modified Gaussian naïve Bayes model) is trained before being used. It should be understood that this part of the training process is carried out, for example, on a general purpose computing device and not on an apparatus such as that illustrated in Figure 1. As shown in Figure 2, at step 200 various feature training data sets are used as inputs. Then for each combination of feature and class, at step 201 full precision constants and coefficients are computed. An iterative process then begins to step through each feature and class to determine the respective class probabilities (see Equation 6 above) for each possible input value. Note that in this example the inputs will be quantized to a 5-bit value (though the present techniques are not limited to this quantization of input value), so there are 32 possible input values for each input. The inventors of the present techniques have found that the class prediction accuracy drops only by around 1% when such quantised input values are used. Hence at the first iteration at step 202 the first feature (Feature₀), at step 203 the first input (Input = 0), and at step 204 the first class (Class₀) are set to be considered. For these parameters, at step 205, the corresponding class probability is determined (see Equation 6 above). It is then checked at step 206 if all classes (for this feature and input value combination) have now been considered. In some embodiments described here there are 8 classes (though the present techniques are not limited to this particular number of classes). Whilst they have not, step 207 gets the next class to consider and the flow returns to step 205. Once all class probabilities for all classes (for this feature and input combination) have been determined the flow proceeds to step 208, where the class with the maximum probability is found. This class ("MaxClass") is then stored into class indications storage, such as a lookup table for this feature and input value combination at step 209. At step 210 it is then determined if the last input value (in this 5-bit example, this being number 31) has been reached. Whilst it has not step 211 increments the input value and the flow returns to step 204. Once the last input value has been reached the flow proceeds to step 212, where it is determined if the last feature of the set has been reached. Whilst it has not step 213 gets the next feature to consider and the flow returns to step 203. Once the last features has been reached,

then the full set of iterations (over features, input values and classes) is complete the flow of the training process concludes at step 214.

5 An algorithm is used by the present techniques, according to which instead of considering a single naïve Bayes classifier with d features, d distinct Bayes classifiers are considered and their predictions are aggregated to select the final classification. A useful characteristic is that each feature (represented by a feature value) may derive from a completely different distribution (e.g. Gaussian, exponential, uniform, etc.).

10 Figure 3A shows an example set of distributions according to which each of four feature value data sets are assumed to be represented by Gaussian distributions, whilst Figure 3B shows a set of four feature value distributions in which two are Gaussian distributions and two are uniform distributions.

15 Figure 4 schematically illustrates an apparatus 400 in some example embodiments. Here the class retrieval circuitry 401 comprises 4 lookup tables (LUT) 402, one of which is provided for each feature data value which the apparatus is configured to receive – note that in this example the feature dataset input circuitry is not explicitly represented. Accordingly, on receipt of a feature data set, each feature
20 data value is used in a lookup in a respective LUT 402 and a class indication from each is read out. This set of class indications is then passed to class selection circuitry 304 which selects a single representative class on the basis of the class indications received.

25 Figure 5 schematically illustrates an apparatus 500 in some example embodiments. Feature data values are received by a feature data input data 501 which holds these values such that the lookup control circuitry 502 can make use of them in sequence to perform a lookup in the single lookup table 503. That is to say four lookups in the lookup table 503 are performed in a serial sequence under the control of
30 the lookup control 502 using the four feature data values received. The result of each lookup in the lookup table 503 is passed to the class determination circuitry 504 where, whilst the serial lookup procedure is being carried out, the classes received are

temporarily stored in class storage 505. From here they are then retrieved for the class voting circuitry 506 to determine a single class for output on the basis of a vote between them. Here the vote is performed by majority in the sense that the most common (mode) class is selected as the winning class for output.

5

Figure 6A schematically illustrates one example embodiment in which weights 600-603 are associated with the classes retrieved from the class indication storage before a class selection takes place. These weights may be learnt as part of the training, as for example carried out as shown in Figure 2, or may be explicitly set by a user, wishing to influence the balance of the class selection. Hence, for each class received from the class indications storage (e.g. a lookup table) an associated weight is used for each. The weights can be stored in the same lookup table as the classes or they can be in a separate storage. It is to be noted therefore that this weight is thus effectively applied for the “importance” of each respective feature data value, but the class values themselves are not amended since these are integer values used for enumerating the set of possible classes. Each class and its associated weight are then received by the class selection circuitry 604 which performs a determination of the selected class on the basis of the distribution of classes themselves, and their associated weights. For example, the weights may indicate the relative voting weight that each feature data value thus has in the final selection. Additionally, or as an alternative, the class selection circuitry 604 may comprise tie breaker circuitry 605 which can make use of the weight in the event that a tie break is necessary. A tie break occurs when the selection procedure (e.g. mode voting) cannot differentiate between two or more selected classes. In this instance the respective associated weights may be used as a tie break influence. As mentioned above the selection between the class indications received from the class indication storage may take a variety of forms, but in some embodiments such as that illustrated in Figure 6B, this selection is by vote, for example by mode vote.

30

Figure 7 schematically illustrates some components of an apparatus 700 in one embodiment. In this example eight feature data values are received, each quantised into 5-bit values. These respective 5-bit values are used to perform lookups in a

corresponding set of eight lookup tables (of which only two 701 and 702 have been explicitly shown, purely of the sake of clarity) which have each been pre-populated (in a training phase) with 3-bit class indications for each possible value of each feature value. In other words, there are 32 entries in each lookup table. Hence, eight lookup
5 actions are performed in parallel to retrieve eight 3-bit class values which are passed to the class voting circuitry 703. The class voting circuitry 703 then performs the final selection of the class prediction which in this example is on the basis of a mode vote amongst the eight class indications received. An implementation such as that shown in Figure 7 (using the new algorithm described herein) has been demonstrated to have a
10 class prediction accuracy of 91%, which is very close to the accuracy of the Gaussian naïve Bayes.

The present techniques may find implementation in variety of contexts, but Figure 8A gives an example of particular implementation where the feature data set is
15 generated by a set of sensors. Accordingly, an apparatus 800 is shown which receives respective feature data values from four sensors 801-803. These are received by the feature data set input 804 which passes them to the class lookup table 805 in order for a set of four class indications to be retrieved. The class voting circuitry 806 (as described above) then selects a single one of these class indications as the final class
20 output, here by mode voting. As in the example of Figure 8A the sensors may be external to the apparatus but other examples are possible, one of which is illustrated in Figure 8B. Here the apparatus 810 is a self-contained unit wherein four sensors 811-814 form part of the apparatus 810. As in the example of Figure 8A the output of these sensors is received by a feature dataset input 815 which can temporarily hold
25 these values before they are passed to a set of class lookup tables 816, in order for a corresponding set of class indications to be read out. These class indications are passed to the class voting circuitry 817 which then determines a single class on the basis of a vote between them. The apparatus 810 in this example further indicates four indicators 818-821 which are used by the apparatus to indicate which class was
30 selected. Accordingly it will be understood that four different classes are defined in this example, but note that the present techniques are not constrained to this number of classes (and further that the fact that four inputs are shown from four sensors is purely

coincidental). These indicators may take a variety of forms, but to name just one example may be a visual indication, e.g. illuminating a light or changing the colour of a small surface area or a different LED for each class, such that a user can perceive that on the basis of the current sensor data input a particular class has been selected.

5 Accordingly, in such an example the class indications can characterise different situations which may be linked to different balances of sensor data input. It is to be further noted that the apparatus 810 may be physically constructed in a variety of different ways. This could for example be a small system on chip embodied in silicon but in other examples (in particular such as will be described below with reference to
10 Figure 9, the apparatus may instead be embodied as a plastic fabricated device. Although according to contemporary technologies the number of logic gates which can be provided for a given area of plastic fabricated device by comparison to a silicon fabricated device is considerably lower, the particularly low gate counts which are possible according to the present techniques have been found to lend themselves
15 particularly well to such plastic fabricated devices. The above discussed implementation shown in Figure 7 has been implemented in hardware using such plastic technology. The class prediction can be made in $10\mu\text{s}$, which is 2-8x faster than a Gaussian naïve Bayes implementation. Moreover it consumes only 315 gates, which is 10-30x smaller than a Gaussian naïve Bayes implementation.

20

Figure 9 shows the torso area of a user in an example in which the apparatus may be a wearable device. Hence in this example the apparatus 901, which may for example be arranged in accordance with the example of Figure 8B as described above, is worn on or close to the skin of a user 900. In just one example embodiment the
25 sensors which form part of the apparatus are then configured to be responsive to local environment conditions which it is desirable to monitor. These could be variously configured from a range of known sensors. Any available sensors may be deployed here, but examples such as sensors for temperature, humidity, pressure, ECG/EMG, or the presence of particular chemicals may be contemplated. It should be noted that the
30 position of the apparatus in the example of Figure 9 is merely for clarity of illustration and this device could be worn in any appropriate location on or near to the skin. For example in one contemplated example the apparatus embodied as a plastic fabricated

device might be worn underneath the armpit, such that a range of chemical sensors could determine the balance of chemicals present and activate one of the indicators to signal to the user a particular message about the current “chemical balance” of the under arm area.

5

Figure 10 shows a sequence of steps which are taken accord to the method of one example embodiment. At step 1000 a feature data set is received and then at step 1001 each bit set (representing each feature value in the feature data set) is used to perform a lookup in a lookup table of class indications. The corresponding set of class
10 indications is then read out at step 1002 and at step 1003 a vote is conducted based on the set of class indications read out to determine a selected class. The selected class is then output at step 1004.

In brief overall summary, apparatuses and methods of operating such
15 apparatuses are disclosed. An apparatus comprises feature dataset input circuitry to receive a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits. Class retrieval circuitry is responsive to reception of the feature dataset from the feature dataset input circuitry to retrieve from class indications storage a class indication for each feature
20 data value received in the feature dataset, wherein class indications are predetermined and stored in the class indications storage for each permutation of the set of bits for each feature. Classification output circuitry is responsive to reception of class indications from the class retrieval circuitry to determine a classification in dependence on the class indications. A predicated class may thus be accurately
25 generated from a simple apparatus.

In the present application, the words “configured to...” are used to mean that an element of an apparatus has a configuration able to carry out the defined operation. In this context, a “configuration” means an arrangement or manner of interconnection
30 of hardware or software. For example, the apparatus may have dedicated hardware which provides the defined operation, or a processor or other processing device may be programmed to perform the function. “Configured to” does not imply that the

apparatus element needs to be changed in any way in order to provide the defined operation.

Although illustrative embodiments have been described in detail herein with
5 reference to the accompanying drawings, it is to be understood that the invention is not
limited to those precise embodiments, and that various changes, additions and
modifications can be effected therein by one skilled in the art without departing from
the scope of the invention as defined by the appended claims. For example, various
combinations of the features of the dependent claims could be made with the features
10 of the independent claims without departing from the scope of the present invention.

CLAIMS

1. Apparatus comprising:
feature dataset input circuitry to receive a feature dataset comprising multiple
5 feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits;
class retrieval circuitry responsive to reception of the feature dataset from the feature dataset input circuitry to retrieve from class indications storage a class indication for each feature data value received in the feature dataset, wherein class
10 indications are predetermined and stored in the class indications storage for each permutation of the set of bits for each feature; and
classification output circuitry responsive to reception of class indications from the class retrieval circuitry to determine a classification in dependence on the class indications.
15
2. The apparatus as claimed in claim 1, wherein the class indications stored in the class indications storage for each feature are each predetermined as a best class indication which maximises a Bayes Classifier for the feature in a training phase using feature training datasets.
20
3. The apparatus as claimed in claim 2, wherein the Bayes Classifier is a Gaussian naïve Bayes Classifier.
4. The apparatus as claimed in claim 2 or claim 3, wherein the Bayes Classifier is
25 based on a single distribution type used for each feature of the set of features.
5. The apparatus as claimed in claim 2 or claim 3, wherein the Bayes Classifier is based on heterogeneous distribution types used for the set of features.
- 30 6. The apparatus as claimed in any preceding claim, wherein the class indications storage has a look-up table format and the class retrieval circuitry is arranged to

perform a look-up procedure with respect to the look-up table format for each feature data value received in the feature dataset.

7. The apparatus as claimed in any preceding claim, wherein the class retrieval
5 circuitry is arranged to retrieve in parallel the class indications for each feature data value received in the feature dataset.

8. The apparatus as claimed in any of claims 1-6, wherein the class retrieval
10 circuitry is arranged to retrieve in a serial sequence the class indications for each feature data value received in the feature dataset.

9. The apparatus as claimed in any preceding claim, wherein the classification
output circuitry is responsive to reception of the class indications from the class
retrieval circuitry to determine the classification by a vote amongst the class
15 indications.

10. The apparatus as claimed in any preceding claim, wherein the class indications
are weighted.

20 11. The apparatus as claimed in claim 10, wherein weights for the class indications are determined in a training phase using feature training datasets.

12. The apparatus as claimed in claim 10 when dependent on claim 9, wherein
weights of the class indications are used as a tiebreaker when the vote selects more
25 than one class indication.

13. The apparatus as claimed in any preceding claim, wherein each feature data
value is represented by a set of fewer than 10 bits, preferably by a set of 5 bits.

30 14. The apparatus as claimed in any preceding claim, wherein the class indications are stored in the class indications storage using a representation of fewer than 5 bits, preferably using a representation of 3 bits.

15. The apparatus as claimed in any preceding claim, wherein the apparatus is a wearable device.

5 16. The apparatus as claimed in any preceding claim, wherein the feature dataset input is coupled to a plurality of sensors each providing a respective feature data value.

17. The apparatus as claimed in any preceding claim, wherein the apparatus is embodied as a plastic fabricated device.

10

18. The apparatus as claimed in any preceding claim, wherein the apparatus is embodied comprising fewer than 1000 logic gates and preferably comprising fewer than 500 logic gates.

15 19. A method of operating an apparatus comprising:
receiving at a feature dataset input a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits;
retrieving from class indications storage a class indication for each feature data
20 value received in the feature dataset, wherein class indications are predetermined and stored in the class indications storage for each permutation of the set of bits for each feature; and
determining a classification in dependence on the class indications.

25 20. Apparatus comprising:
means for receiving a feature dataset comprising multiple feature data values indicative of a set of features, wherein each feature data value is represented by a set of bits;
means for retrieving from means for storing class indications a class indication
30 for each feature data value received in the feature dataset, wherein class indications are predetermined and stored in the means for storing class indications for each permutation of the set of bits for each feature; and

means for determining a classification in dependence on the class indications.

5

10

1/10




FIG. 1

2/10




FIG. 2

3/10




FIG. 3A




FIG. 3B

4/10




FIG. 4

5/10




FIG. 5

6/10




FIG. 6A




FIG. 6B

7/10




FIG. 7

8/10




FIG. 8A




FIG. 8B

9/10




FIG. 9

10/10




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2020/050469

A. CLASSIFICATION OF SUBJECT MATTER
INV. G16H50/20
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G16H

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Anonymous: "ASCII - Wikipedia", 7 February 2019 (2019-02-07), XP055689521, Retrieved from the Internet: URL:https://en.wikipedia.org/w/index.php?title=ASCII&oldid=882225814 [retrieved on 2020-04-27] page 1 par 1, page 2 par 1 -----	1-20
A	Anonymous: "Program to implement ASCII lookup table - GeeksforGeeks", 11 December 2015 (2015-12-11), XP055689525, Retrieved from the Internet: URL:https://www.geeksforgeeks.org/program- implement-ascii-lookup-table/ [retrieved on 2020-04-27] page 1 par 1 -----	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

11 May 2020

Date of mailing of the international search report

25/05/2020

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Bankwitz, Robert