

Description

The present invention relates to pattern recognition, and particularly, though not exclusively, to speech recognition.

- 5 The n-tuple method of pattern recognition; originally suggested by Bledsoe & Browning ("Pattern recognition and reading by machine", Proc. Eastern Joint Computer Conf., Boston, pp225 - 232; 1959) has been proposed for the recognition of two-dimensional patterns. Figure 1 shows an NxM pattern each element of which is represented by a '0' or a '1'; a single bit. Sets of n bits are selected in a specified way (e.g. at random) from the array forming in each case an 'n-tuple'. Usually each bit is used once so that there are NM/n n-tuples.
- 10 A template store (figure 2, which assumes n=4) has NM/n rows (one for each n-tuple) and 2^n columns. In a training sequence each n-tuple is interpreted as a binary number from 0 to 2^n-1 and a 1 is written into the corresponding column of the row assigned to that n-tuple. A number of training passes on patterns all of course with the same n-tuple selection will plot further 1's into the template store - which may or may not coincide with those already written in, according to the degree of similarity between the patterns. As only 1s and 0s are recorded, single occurrences and multiple occurrences are indistinguishable, each being marked by a single 1.

- Templates are formed in this way for a number of patterns to be recognised. When an unknown pattern is to be identified, n-tuples are formed in the same way and each is used to read out the corresponding location in one of the template stores. The number of '1's found represents a 'score' of the similarity between the unknown and the known pattern. Scores are obtained for each template and the unknown pattern is deemed to be recognised as being that corresponding to the template giving the highest score.
- 20 Tattersall and Johnston ("Speech Recognisers based on N-tuple Sampling", Autumn conference, 1984 Proceedings of the Institute of Acoustics, Vol 6 part 4 pp405 - 412) have proposed a speech recogniser using this principle. In this case in the pattern of figure 1 the column represents successive samples in time of a word of speech, and the bits within each column represent a binary coding (e.g. a bar code) of that sample (or optionally of extracted features). This pattern is then analysed in much the same way as described above.

- 25 In US patent 4087630 there is described an apparatus for speech recognition in which a speech signal is sampled, digitized and reduced to a string (or sonogram) of binary digits each of which digits indicates the presence or absence of significant amplitude in one of a plurality of frequency bands used in sampling for the relevant time interval. Sonograms for several time slices are then concatenated, randomized and then decoded using an n-tuple technique to produce a pattern corresponding to the then current speech signal. This pattern is learnt by superimposing it on an existing vocabulary entry and is subsequently recognized if it is sufficiently similar to one vocabulary entry and different from all others.

- 30 In US 4087630 there is described an apparatus for speech recognition in which a speech signal is sampled, digitized and reduced to a string (or sonogram) of binary digits each of which digits indicates the presence or absence of significant amplitude in one of a plurality of frequency bands used in sampling for the relevant time interval. Sonograms for several time slices are then concatenated, randomized and then decoded using an n-tuple technique to produce a pattern corresponding to the then current speech signal. This pattern is learnt by superimposing it on an existing vocabulary entry and is subsequently recognized if it is sufficiently similar to one vocabulary entry and different from all others.

- 35 US 4087630 in the training mode neither records for each of the n-tuples the frequency with which combinations of bits within the n-tuple occur in a plurality of patterns, nor records for each n-tuple data indicating which combinations of bits within that n-tuple occur in the training patterns. Similarly, stored frequencies are not used during recognition. In addition there is of course no adjustment of any stored frequency.

- 40 In US 3,521,235 there is described a pattern recognition technique in which, in the learning phase, frequencies of occurrence of selected binary words occurring within digital codes derived from patterns of the type to be identified are used to determine which binary words best identify and distinguish the codes and hence the patterns. The "distinguishing" words which have been selected in this way are used subsequently in the recognition process in preference to other words which are less capable of distinguishing between the patterns. The frequencies of occurrence obtained in the training process are assigned weighting factors to facilitate discrimination between different classes of patterns. The assigned weighting factors rather than the frequencies of occurrence are used in the recognition process.
- 45

- 50 In "Speech recognition using autocorrelation analysis", by R.F. Purton, IEEE Trans. Audio and Electroacoustics, Vol. AU-16, No. 2, 1968, pp235 - 239, there is described a speech recognition technique in which a master pattern or template is produced for each word of the chosen vocabulary by combining the individual binary patterns from several known utterances of the word. The combining process is a simple cell-by-cell addition, so that the number in any cell of the master pattern is equal to the number of individual patterns in the reference or training group which contained a 1 in the corresponding cell; the number is thus proportional to the probability of a 1 occurring in that position. A further optional facility enables a master pattern to be modified or updated by the addition of a correctly recognized pattern.

- 55 The present invention is defined in the appended claims.

Some embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings in which:

Figure 3 is a block diagram of one form of speech recognition apparatus according to the invention;

Figure 4 illustrates the operation of the template store of figure 3;

Figure 5 illustrates a modified form of microfeature selection; and

Figure 6 illustrates a modified version of the template store of Figure 4.

The embodiment to be described is concerned with speech recognition; however, it is to be noted that the invention is also of value in recognising other types of pattern.

5 The apparatus shown in figure 3 has an audio input 10 for receiving speech signals which are converted by an analogue to digital converter 11. Although the raw time domain samples could be directly utilised (as proposed by Tattersall and Johnston), in this embodiment they are supplied to a front end processing unit 12 which serves to convert them (as described in more detail below) into binary representations of various characteristics of the speech e.g. energy profile, spectral values, zero-crossing rate. It is assumed that this output
10 is of 80 bits and is provided for each 10ms period of speech.

The front end processing unit 12 supplies data continuously, to an input array store 13, but the contents of the store are not acted upon until the commencement of a word is recognised by a detector 14 which monitors the input signal and triggers when for example a certain zero crossing count or energy threshold - or a combination of the two - occurs. A few periods of data - eg 30ms - are also captured to ensure that low level sounds
15 occurring before the trigger event occurs are included; a delay 15 may be included if the inherent delay of the front end processing unit 12 is not sufficient.

The input array store 13 accommodates 1 second of data (ie 100 10ms periods) and therefore has a capacity of $80 \times 100 = 8000$ bits. However the detector 14 is arranged to inhibit the further processing of periods of silence in the event that the word input is shorter than 1 second duration. If desired, the duration of the word
20 may be taken into account in weighting the final recognition process.

The 8000 bits are notionally divided into microfeatures each consisting of an ordered set of 8 bits ("8-tuples"). In principle these are selected at random from the input array store although as will be discussed further below certain constraints may be placed on the selection. In the case of figure 3 there are 1000 microfeatures. The addresses of those bits in the input array making up each microfeature are stored in an assignment register
25 16 which is a store having 1000×8 locations accessible by means of an index number (1...1000) and bit number (1...8) produced by an input control unit 17.

Note that it is not in principle necessary that all microfeatures consist of the same number of bits; however this will be assumed in the description which follows.

The apparatus operates in (a) a learning mode and (b) a recognition mode. In the learning mode, spoken
30 words input in a training sequence are used to form templates in a number of template stores 18. Unlike the previous proposals discussed above which simply record the fact of occurrence of a given microfeature having a particular value, it is proposed here to record the frequency of its occurrence during the training sequence. Although in principle a similar store organisation might be used, but with several bits per location to accommodate a frequency rather than just a single bit, this represents a significant storage capacity (viz. 1000×256
35 $\times$ eg 3 bits for the frequency = 768Kbit) per store, and the preferred arrangement, which also serves to facilitate certain modifications, is as described below. Note however that this store arrangement is advantageous even in the case where the frequencies are not stored.

One template store 18 is depicted in figure 4 (the others being identical). It contains several (eg 5) entries for each microfeature (ie 5000 entries in all); each entry consists of one 8-bit byte for a microfeature value,
40 and one byte for its frequency (ie 10,000 bytes).

The index number determines the address where the microfeatures are stored, and provided the number of locations reserved for values of each microfeature is fixed, then the location and index are functionally related, and each can readily be calculated from the other.

During the training sequence, a chosen word is repeated a certain number of times (eg 5). Successive
45 microfeatures are generated by applying the index numbers (and bit numbers) in turn to the assignment register 17 to address the appropriate location in the input array store 13, and provide 1000 8-bit microfeature values. Each microfeature value obtained is stored as an entry in the template store with a frequency of one, unless it has previously occurred for that microfeature in which case the frequency stored against that value is increased by one. Thus if, as shown for microfeature index no 738, value 201 has occurred four times and 137
50 once, there are two entries for that microfeature, the other three being zero.

In the recognition mode, microfeature values are obtained from the input array store in exactly the same way. Referring again to figure 3, a search unit 19 takes each value and searches in the template store the group of entries corresponding to that microfeature to find a match; if a match is found, the corresponding frequency is read out and passed to an accumulator 20 which forms the sum of the frequencies thus read out. This result
55 is then divided (divider 21) by the maximum possible total for that template, ie the sum of the maximum frequency for each of the 1000 microfeatures. This total could be calculated when required, or (as shown) may be stored in an additional location in the template store.

A template store 18, accumulator 20, and divider 21 are provided for each word to be recognised. A com-

parator 22 receives the scores output from each of the dividers 21 and produces at an output 23 a signal indicating which template has produced the highest score - ie which word has been recognised. The recognition process may be subject to the criterion that the highest score exceed a threshold before any output occurs.

5 The apparatus described relies for time alignment upon the action of the detector 14. It is possible however that variations in volume (for example) may result in inconsistent alignment of the training utterances. This means that the apparatus having been trained on several utterances with differing alignments 'learns' a degree of tolerance to misalignment. However, in a modified version of the apparatus, each word after the first is subjected to the recognition process against the template already found, both in the time alignment determined by the detector 14 and also with shifts of (for example) $\pm 10\text{ms}$ and $\pm 20\text{ms}$. The alignment which gives
10 the highest score is chosen and the template updated on that basis. This adjustment will reduce the above-mentioned tolerance of the recognition process to time alignment errors and therefore it follows that the same trial shifting of the unknown utterances must be carried out prior to recognition. The desirability of this modification will depend on the time delay acceptable for recognition where the template store is to be trained on different words; the alignment process would of course need to operate independently for the different words.
15 A further preferred feature is the provision of adaption, whereby the templates continue to "learn" during the recognition mode. A simple method of achieving this would be to increment the frequency associated with a microfeature value which contributes to a successful recognition. Thus the "frequencies" would become the frequencies of recognitions known to be successful, rather than the frequencies of occurrence during the training session.

20 The increasing of the frequencies by one could be done automatically, whenever the recogniser operates, on the assumption that the word output is always correct, but preferably occurs in accordance with an external input so that the machine is told when it is right, or (alternatively) only if it is wrong. If the machine is used in a dialogue, then the dialogue might be configured so that confirmation is requested at various stages. When this happens in a way which indicates a successful recognition, the frequencies of the microfeature values
25 which contributed to the score would be increased by one.

A more sophisticated technique would be to add to the template all the microfeatures from the successfully recognised word. This facility would require the provision of increased storage capacity for the template, and a strategy for preventing the stored frequencies becoming too large. The corollary of this is that the template can be allowed to "forget" microfeature values which no longer appear. This implies that the apparatus when
30 used for example for speaker recognition or verification will be able to track changes in the speaker's voice over a period of time.

Figure 5 shows a section of a modified template store; only the storage allocated for one microfeature (eg index no. 312) is shown, but each microfeature is allocated 16 entries of two (8-bit) bytes each. During the learning phase (eg) five utterances yield up to five values for microfeature 312, with frequencies totalling 5.
35 Typical figures are shown in figure 5a, in descending order of frequency, in the first three entries. The remaining entries have a frequency of 0; the contents of the associated value fields are of course immaterial.

In the recognition mode, each time the template gives rise to a successful recognition, either:

(a) the microfeature has a value not already entered; this value is entered in the next available entry field with a frequency of one; (if - as described below for time warping - more than one value is obtained, the
40 choice of which one to enter is arbitrary), or
(b) the microfeature has a value already entered; the existing frequency is increased by one and (if necessary) the order of the entries is adjusted to maintain the descending order of frequency. In this way, the storage reserved for each microfeature will after a time consist of several microfeature values at the top of the list with frequencies greater than one, and the remainder will have frequencies of one. (Figure 5b).

45 However it is suggested that only the first eight entries are used in the recognition process; the values in locations 9 to 16 being kept in waiting, and becoming effective only when the re-ordering process brings them up into the top 8 locations. To avoid values in the top 8 locations varying too rapidly, a notional barrier may be drawn between locations 8 and 9; a value in location 9 can only move up across this barrier if its frequency is more than twice the frequency of the value in location 8.

50 Two possible overflow conditions may occur. Firstly, the upper entries may become too large - in which case all the frequencies are halved, being rounded down so that frequencies on one become zero, ie leaving empty entries for new values to occupy. To prevent the lower locations being filled by ineffective values, this may be done at a much lower frequency (eg 31) than that dictated by the store size. The likelihood of the second possibility - that of all entries being occupied - will thereby be much reduced but if it occurs it is suggested
55 that an existing entry in entry 16 is always jettisoned in favour of a new one.

It should be noted that after halving the "frequencies" of existing microfeature values, and creating vacancies, any new microfeature value will acquire a greater frequency relative to these than its true frequency of occurrence would suggest. This gives it a greater chance of surviving and rising towards the top of the list, a

characteristic which appears to be desirable. In this way, the constitution of the template is continually adapted according to the success the apparatus achieves in recognition. In speaker verification applications, it would clearly be important to adapt the template only when there was 100% certainty that the correct speaker had been identified. Conversely, the confidence that this was so should increase continually the more the apparatus was used.

It has already been mentioned that the selection of microfeatures may not be entirely random. In a modification it is proposed that the ordered set of bits making up a microfeature be randomly selected subject to the constraint that the bits making up any given set all lie within a time window of set width. For example using the figures given above where the input array store has 100 columns each representing 10ms the 8 bits forming each microfeature may lie with a 50 ms window. For example microfeature 1 shown in figure 6 lies within 10ms periods 1 to 5 whilst microfeature 2 lies within 10ms periods 2 to 6 and 3 within periods 3 to 7. The 50ms windows may be arbitrary, or each microfeature may lie within a respective different 50ms window (subject to adjustment at the ends of the store). A second constraint is related to the meaning of the bits in each column of the input data store. As will be discussed below it is envisaged that the 80 bits may consist of several groups each of which represents a particular parameter of the input speech. For example two or three bits may represent a voiced/unvoiced measure either in a binary representation or in a barchart representation:

barchart	binary	
000 -	00	indeterminate (signal level too low)
001 -	01	unvoiced
011 -	10	intermediate voicing
111 -	11	voiced

In the barchart representation, the significance of each bit is independent of the others (ie middle bit = 1 means that at least some voicing is present) and represents no constraint on the microfeature selection. In the case of the binary representation this is not so and the pair of bits is treated as indissoluble; ie the two bits always form part of the same microfeature. This will require some rule to prevent the microfeature having more than 8 bits (this rule will destroy the strict randomness of the choice of bits, but not so as to imperil the technique).

It is to be noted that the time-window feature facilitates the termination of the recognition process if a word ends before the one second period has expired. If random or other arrangements are used, then preferably any microfeatures consisting entirely of bits derived from periods of silence are ignored in the recognition process.

A further proposal is the introduction of time warping, to accommodate variations in the speed at which a given word is spoken. If a word is spoken slowly compared with the words on which the template has been trained, then though microfeatures at the start of the word may coincide temporally, those in the middle or at the end of the word will occur late.

This eventuality may be accommodated by applying the time warping technique now to be described. It should be noted that this assumes accurate temporal alignment of the beginning of the word. Since the time alignment process as described above tends to optimise the average alignment over the whole word that process, if used, is preferably modified by restricting the comparison to the first (eg) 200 ms of the word. As with time alignment, the time-warping technique applied to template training will reduce the template's tolerance to variations in the speed of utterance unless the technique is also applied in the recognition mode.

Assuming firstly that the basic microfeature values have been assigned on the basis that each spans a maximum of 50 ms, extra microfeatures are generated as follows:

The earliest 200 microfeatures generate only one value. The next 200 will generate two values, with the bits coming from two windows displaced in time by 20 ms but having identical assignments within those windows. (If the assignment operator is visualised as an eight-pronged fork, highly distorted, stabbed into the data array, to pick out the bits forming a microfeature, then the time-warped microfeature value is generated by stabbing the same fork on to the data array 20ms later).

The next 200 will generate three values, from three windows each displaced in time by 20ms, and so on up to the last 200, which will exist in five forms. The total number of microfeatures will be $200 \times (1+2+3+4+5)$, ie

3000, under this scheme. The spread in time of the last 200 will be 100ms, representing a $\pm 5\%$ time warping. The actual degree of time warping applied would need to be chosen on the basis of experiment; the figures quoted being merely by way of example.

During training then microfeature 201 (for instance) will have two values for each utterance; assuming five
 5 training utterances there will be ten values. If there is no duplication of a value across the different utterances, then one value is chosen arbitrarily from each utterance and given a frequency of 1. Each value is given a frequency (ie the number of different utterances which produced that value) and the five values with the highest frequency are stored in the template. Similarly, for the later microfeatures with multiple values: each value is
 10 given a frequency equal to the number of different utterances producing the same value, and the five values with the highest frequencies are stored in the template for the chosen word.

In recognition mode, the information stored in the input array store is processed to produce 3000 micro-
 feature values. The values of these are used to form a data array, referred to here as the "tableau", which has
 15 dimensions of 1000 by 256 and forms part of the search control 19 of figure 3. The index (from 1 to 1000) of each microfeature addresses a row of this tableau, and a 1 is written into the column corresponding to the value (from 1 to 256) of the microfeature. Thus the first 200 rows will have one 1, rows 201 to 400 one or two 1's,
 and so on with up to five 1's in the last 200 rows.

The tableau is now addressed by the stored microfeature values forming the template for each word. Spe-
 cifically, each microfeature value from a template addresses the row in the tableau having the same index,
 20 and the column with its value, and if a 1 is present, then the frequency of the microfeature value is treated as its weight, read out, and accumulated to arrive at a total score. As before, this is then normalised by dividing by the maximum score the unknown word could achieve, to arrive at the relative probability that the unknown word matches the template word.

For microfeatures 201 onwards, there are two or more entries (1's) in each row of the tableau. Preferably
 25 the microfeature values in the template store are ordered with the highest frequencies at the top, so that the first frequency to be read out will be the highest possible score for that microfeature. As soon as a match has been found the next microfeature is addressed. Similarly, in determining the maximum score for any template, only the frequency at the top of the list corresponding to each index needs to be read out and accumulated.

This is repeated for all the stored templates, and the identity of the template achieving the highest relative
 probability is output as the word deemed to be the best match to the unknown word.

30 This warping process represents a quasi-linear distortion of the time axis and is easy to do by simple shift-
 ing because of the restricted time window within which the bits of any given microfeature lie. If entirely random assignment of the bits is used then the individual bits of the microfeature would need to be shifted by amounts depending on their temporal position within the input data array (analogous to an elastic fork). Those micro-
 35 features which happened to contain bits in the last 200ms would exist in 5 forms, with the individual bits suitably warped in a quasi-linear manner. Those which contained bits no later in time than the interval 600 to 800ms from the start of the input data array would exist in 4 forms, and so on, so that a microfeature all of whose component bits fell within the initial 200ms interval would exist in only one form.

Statistically, the total number of microfeature values will be just the same under this scheme, ie 3000. The
 operation of this technique will be very similar to the technique described above.

40 The nature of the front end processing 12 will now be discussed further. There are a number of possible
 parameters which may be derived for entry into the input data store. Not all of these would necessarily be used in any particular apparatus, though since it is envisaged that the parameters will be fairly coarsely quantised it is thought desirable to include both frequency domain parameters and time domain parameters, on the basis
 that the information lost in quantisation is different in the two cases. Possible features are as follows:

- 45 i) total energy within the speech frequency band
- ii) voiced/unvoiced measure
- iii) spectral energy levels - ie. the energy levels within a number of respective frequency bands spanning the range.
- iv) spectral energy gradients - ie. the difference between the levels in adjacent ones of the bands referred to in (iii).
- 50 v) spectral energy time profiles - ie. the differences between the energy levels within a number of respec-
 tive frequency bands in successive sample periods
- vi) zero-crossing rate - the number of zero-crossings of the speech wave form during a given period
- vii) wiggle rate - the number of stationary points (maximum and minimum) in the speech wave form during
 55 a given period
- viii) excess stationary points - the difference between (vii) and (vi).

Considering possible implementation for each of these in turn, assuming a sampling rate of for example
 one per 10ms:

i) Energy

A filter spanning the whole of the useful speech frequency range, say 200Hz to 5000Hz (6db points, gentle roll-offs), is used to produce an energy profile of the word. The absolute output level of the filter is represented by four bits, ie 16 levels which are then coarsely quantised, with non-linear intervals; ie each of the 16 levels will be assigned to one of (say) four or five new measures. For a continuous quantity such as energy, bar-chart coding is used, and three possibilities are:

10	00	000	0000	"silence" -	ie energy below some threshold
		001	0001	low	
15	01		0011	moderate	
		011	0111	considerable	
	11	111	1111	loud	

Normalisation could be introduced, in principle, by recording the actual dynamic range of the input energy levels over some suitable period of time, and adjusting the assignment to the parameter codes accordingly. However, agc in speech needs to be treated with caution, and should only act rather slowly.

ii) Voiced - Unvoiced Measure

The ratio of low-frequency energy to high-frequency energy is used as a measure of voicing. A two-bit code should be adequate, ie:

00	indeterminate, ie levels too low for ratio to be significant
01	unvoiced - ratio less than some threshold
10	intermediate voicing
11	voiced

Suitable frequency ranges for the two filters are: low-band, 300Hz, to 1600Hz, 3db points; high-band, 2300Hz to 5000Hz, 3db points, both with gentle roll-offs. The energy levels in the filters would be represented by four-bit linear codes. The assignment of their ratios to the parameter codes would probably be non-linear, and would in any case depend on one or other being above some appropriate threshold level (different for each filter band).

iii) Spectral Energy Levels

Spectral analysis is carried out by a bank of 16 filters, spanning the range 200Hz to 5000Hz, with centre frequencies logarithmically spaced from 250Hz to 4500Hz, and with 4.5db loss at the overlap points. The outputs from the filters are sampled very 10ms, coded linearly by four bits, and are then heavily quantised for assignment to parameter codes. Three or four non-linear levels are probably sufficient, represented by bar-chart coding (two or three bits), ie

	00	000	below some threshold
	01	001	low
50		011	intermediate
	11	111	high

Again, normalisation could be introduced, as discussed under (i) with the same caveat.

iv) Spectral Energy Gradients

The linearly coded outputs from the bank of filters is examined a second time to provide another set of

parameter codes. These are, in each 10ms period, the energy differences progressing from low frequencies to high. There are 15 of these, and they are then quantised very coarsely, as follows:

- 00 indeterminate - both levels less than some threshold
- 01 "-": difference negative
- 5 10 "+": difference positive
- 11 "=": difference less than some threshold

If the voiced-unvoiced measure indicates partial or full voicing, then a sequence (progressing from low frequencies to high) [+ , -], or [+ , = , -] indicates a maximum, and may be identified with a formant. It is known that formant ratios are important "quasi-invariants" of speech. Thus the use of logarithmically spaced filters allows these ratios to be displayed by simple realigning the code, in such a way as to bring all the lowest formants to the same level.

Imagine that the codes above are written on vertical strips of paper, one for each successive 10ms interval. If the lowest entry, corresponding to the difference between the energy levels from the two lowest filters, is +, then the first formant will be assumed to be above the lowest filter. Slide the strip down until the new entry at the original level is either = or -. Then the first formant will now be aligned with (or slightly below) the original level, which is adopted as a datum line. If the lowest entry is =, then it will be assumed that the first formant lies between the two lowest filter bands, and the strip is left unchanged. If the lowest entry is -, then it will be assumed that the first formant lies below the lowest filter band, but hopefully not too far below, and the strip will again be left unchanged.

In this way, the first formants will all be (roughly) aligned with the datum line. Since the filters are logarithmically spaced, the position of the next maximum, indicated by the sequences [+ , -] or [+ , = , -], will indicate the ratio of the second to the first formant, and the position of the next maximum will indicate the ratio of the third to the first.

The realigning of the spectral energy differences is only carried out when partial or full voicing is detected. Otherwise the entries are left unchanged.

The realignment is intended to facilitate speech recognition. For speaker verification purposes, no realignment is carried out. Thus in a transaction where verification is important at certain stages, and recognition at others, a simple change of algorithm could be made under software control.

30 v) Spectral Energy Time Profiles

Just as the energy gradients against frequency can be categorised, within one 10ms time interval, so the energy gradients in each filter band can be categorised against time. Here there are 16 energy differences, which can be represented by the same coding scheme as in (iv) above. If partial or full voicing is being used to realign the spectral energy gradients, as described in the previous section, then the profiles against time can be subjected to the same transformation, which must be controlled by the spectral energy gradient information.

40 vi) Zero-crossing Rate

The number of zero crossings occurring in a fixed time interval, say 2.5ms, is counted, and coded appropriately eg as follows:

- 0000 no zero crossings
- 0001 1 or 2 crossings
- 45 0011 3, 4 or 5 crossings
- 0111 6 to 11 crossings
- 1111 12 to 20 crossings

The maximum number of zero crossings would be 8 x (period in ms), and the period could be increased or decreased if need be. Some form of smoothing or filtering would be desirable before coding, eg averaging three adjacent values, or applying an fir filter with coefficients 1/4, 1/2, 1/4 for example.

vii) Wiggle Rate

the number of stationary points (maxima and minima) in a fixed period would be counted and coded. The period would be the same as for assessing the zero crossing rate, and a similar coding scheme would be used, perhaps after smoothing, ie

000	0 to 2	extrema; or,	00	0 to 3	extrema
001	3 to 6	"	01	4 to 10	"
011	7 to 12	"			
111	13 to 20	"	11	11 to 20	"

viii) Excess Stationary Points

In seeking a more compact parameter from the waveform, which is hopefully meaningful, it may be worth considering the excess number of stationary points over zero crossings. The parameter would vary from 0 to 20. It could be represented after smoothing as follows:

000	0 to 2	excess	extrema
001	3 to 6	"	"
011	7 to 12	"	"
111	13 to 20	"	"

A possible selection from the above, to make up the 80 bit allocation suggested earlier might be:

a)	Total speech energy	4 bits
b)	Number of stationary points: 4x3 bits, ie	12 "
c)	Spectral energy levels: 16x2 bits, ie	32 "
d)	Spectral energy gradients: 15x2 bits, ie	30 "
e)	Voiced - unvoiced measure:	2 "
Total		80 bits

There is a paradox involved in the method of scoring outlined earlier. Suppose that five utterance are used in forming the template, and that four of them are identical, while the fifth has no microfeatures in common with these four. Then a repetition of the repeated utterance will score 100%, while a repetition of the singleton utterance will score 25%, the maximum it could achieve, although it may have equal validity with the repeated utterance.

One way of monitoring this situation is to count the number of zeros added into the score. When this number becomes rather low, it is an indication of a large number of matches between stored values of microfeatures and values from the unknown utterance, but may nevertheless be associated with a low score. Some kind of weighting could be applied to the actual score in this situation. Alternatively, rather than use the frequency of occurrence as the weight to be accumulated, a monotonic function of the frequency could be used, via a look-up table. An approximation to the square root of the frequency is suggested as a suitable function.

Another possible variation is the provision of accelerated processing. In one such option - where microfeatures are derived from restricted time windows - those derived from an early portion of the unknown utterance can be processed to determine eg those 25% of the templates most likely to score highly. By the time this operation has been completed, further microfeature values will have been generated and these can be used for a final sort of the selected 25% (or a three-stage process could be adopted, depending on the processing speed and number of templates).

Alternatively a random selection of microfeatures may be used in a first sort, followed by one or more further sortings. As a back-up to either of these options the normal processing could be carried out too, to apply a correction to any action taken as a result of the accelerated result, if this proved to be incorrect.



Claims

1. A pattern recognition device comprising:
 - (a) input means for providing a pattern of bits from an input signal;
 - 5 (b) means storing assignment data defining groups of bits selected from the pattern;
 - (c) a template store; and
 - (d) control means operable
 - (i) in a training mode to record in the template store, for each of the defined groups the frequency with which combinations of bits within that group occur in a plurality of the said patterns and operable
 - 10 (ii) in a recognition mode in response to an unknown pattern provided by said input means, to retrieve for each defined group of bits selected from said unknown pattern, in response to the bits of the group, the frequency if any which is recorded in the template store for that particular combination of bits, to form a similarity measure for that unknown pattern as a function of the retrieved frequencies, and to produce an output indicative of successful recognition or otherwise in dependence on whether the similarity measure meets a recognition criterion; and
 - 15 (iii) in the recognition mode, for any group in respect of which there is already stored in the template store a frequency for a particular bit combination occurring in the unknown pattern, and in respect of which the recognition criterion is met, to increase the said stored frequency.
- 20 2. A device according to claim 1 including a plurality of template stores, the control means including comparator means for producing an output identifying that one of the template stores for which the similarity measure meets the recognition criterion.
- 25 3. A device according to claim 1 or claim 2 in which, in the recognition mode, the control means is operable for any group in respect of which there is not already stored in the template store a frequency for the particular bit combination occurring in the unknown pattern, if any, in respect of which the recognition criterion is met to write a frequency into that template store in respect of that group having that combination.
- 30 4. A device according to any one of the preceding claims in which the control means performs the frequency adjustment and frequency entry if any conditionally upon the presence or absence of an external signal supplied to the control means respectively confirming or not confirming successful recognition.
- 35 5. A device according to any one of the preceding claims in which the or each template store has for each group a respective section comprising a plurality of entries each comprising a location for storage of a particular bit combination and a location for storage of the frequency associated with that combination.
- 40 6. A device according to any one of the preceding claims, in which the input means is operable to receive speech signals and produce a pattern of bits representing, for an interval of speech, characteristics of the speech signals in successive time periods of that interval.
- 45 7. A device according to claim 6 in which the device is arranged to disregard groups derived exclusively from parts of the pattern representing silence.
8. A device according to claims 6 or 7 in which the assignment data represent groups of bits selected randomly from the pattern.
9. A device according to claim 6 or 7 in which the assignment data represent groups of bits each of which groups comprises bits selected from within a limited time window of the pattern.
- 50 10. A device according to claim 9 in which the majority of the groups are each selected from different time windows, of equal duration.
- 55 11. A device according to claim 9 or 10 in which the said groups are notionally divided into a first category and successively later categories each consisting of groups derived from later time windows than are the groups of the preceding category, and the control means is operable, at least in the recognition mode, to select, in respect of each group other than those in the first category one or more additional groups each consisting of bits temporally shifted, by an amount which is the same for all bits of that group, relative to the bit assignments of the group to which it is related, the number of additional groups selected and the



range of temporal shifts to which they are subject successively increasing for successive categories.

- 5 12. A device according to claim 8 in which the control means is operable, at least in the recognition mode, to select, in addition to the said bit groups, additional bit groups each related to one of the said bit groups and consisting of bits temporally shifted relative to the bit assignments of the related group by an amount related to the temporal position of that bit within the pattern.
- 10 13. A device according to claim 11 or 12 in which the control means is operable to perform the additional bit group selection process also in the training mode after at least one pattern has been processed in the training mode.
- 15 14. A device according to any one of claims 6 to 13 in which the control means is operable, at least in the recognition mode, to perform a temporal alignment process comprising forming the similarity measure in respect of the pattern with a plurality of temporal alignments and processing that pattern with whichever temporal alignment produces the similarity measure indicating the greatest similarity.
- 20 15. A device according to claim 14 in which the control means is operable to perform the temporal alignment process also in the training mode after at least one pattern has been processed in the training mode.
- 25 16. A device according to claim 14 or 15 when dependent on claim 11, 12 or 13 in which the temporal alignment process is restricted to an early portion of the pattern.
- 30 17. A device according to any one of claims 6 to 16 in which the input means is arranged to produce for each of the said time periods a plurality of bits some of which represent a characteristic of the waveform of the speech signal and others of which represent a characteristic of the frequency spectrum of the signal.
- 35 18. A device according to any one of claims 6 to 17 in which the input means is arranged to produce for each of the said time periods a plurality of bits comprising bits representing two or more of the following characteristics:
 - i) total energy within the speech frequency band
 - ii) voiced/unvoiced measure
 - iii) the energy levels within a number of respective frequency bands
 - iv) the difference between the levels in adjacent ones of the bands referred to in (iii)
 - v) the differences between the energy levels within a number of respective frequency bands in successive sample periods
 - vi) the number of zero-crossings of the speech waveform during a given period
 - vii) the number of stationary points in the speech waveform during a given period.
- 40 19. A device according to claim 18 in which the input means is arranged to produce for each of the said time periods bits representing values corresponding to the differences between the energy levels in adjacent ones of the bands referred to in (iii), the said bands having logarithmically spaced centre frequencies, and in which, for each time period, each set of values, considered in an order corresponding to ascending frequency, is, before entry into the said pattern of bits, shifted if necessary such that the lowest-order value which represents higher frequency band energy level equal to or lower than the lower frequency band energy level is aligned with a datum.
- 45 20. A device according to any one of claims 6 to 19 in which at least one characteristic of the speech signal is represented in the pattern by two or more bits the meanings of which are interdependent, and the selection of the bits forming a group or additional group is such that those two or more bits are selected for the same group.
- 50 21. A device according to any one of the preceding claims in which the similarity measure includes forming the sum of the retrieved frequencies.
- 55 22. A device according to any one of the preceding claims in which the similarity measure includes forming the sum of the square roots of the retrieved frequencies.
23. A device according to any one of the preceding claims in which each group has the same number (n) of bits.



Patentansprüche

1. Mustererkennungsvorrichtung mit
 - (a) einer Eingabeeinrichtung, um aus einem Eingangssignal ein Bitmuster zu erzeugen,
 - (b) einer Einrichtung zum Speichern von Zuordnungsdaten, die Gruppen von aus dem Muster ausgewählten Bits definieren,
 - (c) einem Schablonenspeicher, und
 - (d) einer Steuerungseinrichtung, die
 - (i) im Lernmodus es veranlaßt, für jede der definierten Gruppen die Häufigkeit, mit der Bitkombinationen innerhalb der jeweiligen Gruppe in mehreren der Muster auftreten, im Schablonenspeicher zu speichern, und
 - (ii) im Erkennungsmodus es veranlaßt, auf ein unbekanntes, von der Eingabeeinrichtung geliefertes Signal hin für jede definierte Bitgruppe, die aus dem unbekannten Muster ausgewählt wurde, nach Maßgabe der Bits der Gruppe die Häufigkeit, falls vorhanden, die im Schablonenspeicher für diese besondere Bitkombination gespeichert ist, abzurufen, um ein Ähnlichkeitsmaß für das unbekannte Muster als Funktion der abgerufenen Häufigkeiten zu bilden, und eine Ausgabe zu erzeugen, die eine erfolgreiche Erkennung anzeigt oder die in sonstiger Weise davon abhängt, ob das Ähnlichkeitsmaß ein Erkennungskriterium erfüllt; und
 - (iii) im Erkennungsmodus für jede Gruppe, für die bereits im Schablonenspeicher eine Häufigkeit für eine bestimmte, im unbekannten Muster auftretende Bitkombination gespeichert ist und für die das Erkennungskriterium erfüllt ist, die gespeicherte Häufigkeit zu erhöhen.
2. Vorrichtung nach Anspruch 1, mit mehreren Schablonenspeichern, wobei die Steuerungseinrichtung eine Vergleichseinrichtung aufweist, die eine Ausgabe erzeugt, die denjenigen Schablonenspeicher identifiziert, für den das Ähnlichkeitsmaß das Erkennungskriterium erfüllt.
3. Vorrichtung nach Anspruch 1 oder 2, bei der im Erkennungsmodus die Steuerungseinrichtung für jede Gruppe, für die bisher im Schablonenspeicher noch keine Häufigkeit für die im unbekannten Muster gegebenenfalls auftretende Bitkombination gespeichert ist und für die das Erkennungskriterium erfüllt ist, eine Häufigkeit in diesen Schablonenspeicher bezüglich dieser Gruppe mit dieser Kombination schreibt.
4. Vorrichtung nach einem der vorhergehenden Ansprüche, bei der die Steuerungseinrichtung die Häufigkeitsanpassung und Häufigkeitseintragung, wenn überhaupt, dann durchführt, wenn ein der Steuerungseinrichtung zugeführtes externes Signal vorliegt oder nicht, das jeweils die erfolgreiche Erkennung bestätigt oder nicht bestätigt.
5. Vorrichtung nach einem der vorhergehenden Ansprüche, bei der der oder die Schablonenspeicher für jede Gruppe einen jeweiligen Abschnitt hat/haben, der mehrere Eintragungen aufweist, von denen jede einen Ort zum Speichern einer bestimmten Bitkombination sowie einen Ort zum Speichern der dieser Kombination zugeordneten Häufigkeit aufweist.
6. Vorrichtung nach einem der vorhergehenden Ansprüche, bei der die Eingabeeinrichtung Sprachsignale empfängt und ein Bitmuster erzeugt, das für ein Sprachintervall typische Merkmale des Sprachsignals in aufeinanderfolgenden Zeiträumen dieses Intervalls darstellt.
7. Vorrichtung nach Anspruch 6, bei der die Vorrichtung dazu ausgelegt ist, Gruppen außer acht zu lassen, die ausschließlich von Teilen des Musters hergeleitet wurden, die Ruhe darstellen.
8. Vorrichtung nach Anspruch 6 oder 7, bei der die Zuordnungsdaten Bitgruppen darstellen, die zufällig aus dem Muster gewählt wurden.
9. Vorrichtung nach Anspruch 6 oder 7, bei der die Zuordnungsdaten Bitgruppen darstellen, die alle Bits aufweisen, die aus einem begrenzten Zeitfenster des Musters ausgewählt wurden.
10. Vorrichtung nach Anspruch 9, bei der die Mehrzahl der Gruppen aus unterschiedlichen Zeitfenstern gleicher Dauer gewählt wurden.
11. Vorrichtung nach Anspruch 9 oder 10, bei der die Gruppen in eine erste Kategorie und nachfolgende spätere Kategorien unterteilt sind, von denen jede aus Gruppen besteht, die aus späteren Zeitfenstern ab-



- geleitet sind als die Gruppen der vorhergehenden Kategorie, und wobei die Steuerungseinrichtung zumindest im Erkennungsmodus für jede Gruppe außer denjenigen der ersten Kategorie eine oder mehrere weitere Gruppen auswählt, die jeweils aus Bits bestehen, die in bezug auf die Bitzuordnungen der Gruppe, der sie zugeordnet ist, um einen Betrag zeitverschoben sind, der für alle Bits der Gruppe der gleiche ist, wobei die Anzahl der zusätzlichen ausgewählten Gruppen und der Bereich der Zeitverschiebung, die sie erfahren, allmählich für aufeinanderfolgende Kategorien zunimmt.
- 5
12. Vorrichtung nach Anspruch 8, bei der die Steuerungseinrichtung zumindest im Erkennungsmodus zusätzlich zu den Bitgruppen weitere Bitgruppen auswählt, die jeweils mit einer der Bitgruppen in Beziehung stehen und die aus Bits bestehen, die in bezug auf die Bitzuordnungen der in Beziehung stehenden Gruppe um einen Betrag zeitverschoben sind, der mit der zeitlichen Position des Bits innerhalb des Musters in Beziehung steht.
- 10
13. Vorrichtung nach Anspruch 11 oder 12, bei der die Steuerungseinrichtung den Auswahlvorgang für die zusätzlichen Bitgruppen auch im Lernmodus durchführt, nachdem zumindest ein Muster im Lernmodus verarbeitet wurde.
- 15
14. Vorrichtung nach einem der Ansprüche 6 bis 13, bei der die Steuerungseinrichtung zumindest im Erkennungsmodus einen zeitlichen Ausrichtungsvorgang durchführt, der es umfaßt, das Ähnlichkeitsmaß für das Muster mit mehreren zeitlichen Ausrichtungen zu bilden und das Muster mit derjenigen zeitlichen Ausrichtung zu verarbeiten, das das Ähnlichkeitsmaß ergibt, das die größte Ähnlichkeit anzeigt.
- 20
15. Vorrichtung nach Anspruch 14, bei der die Steuerungseinrichtung den zeitlichen Ausrichtungsvorgang auch im Lernmodus durchführt, nachdem zumindest ein Muster im Lernmodus verarbeitet wurde.
- 25
16. Vorrichtung nach Anspruch 14 oder 15, soweit sie von Anspruch 11, 12 oder 13 abhängen, bei der der zeitliche Ausrichtungsvorgang auf einen frühen Bereich des Musters begrenzt ist.
- 30
17. Vorrichtung nach einem der Ansprüche 6 bis 16, bei der die Eingabeeinrichtung dazu ausgelegt ist, für jede Zeitperiode mehrere Bits zu erzeugen, von denen einige ein typisches Merkmal der Wellenform des Sprachsignals darstellen, und von denen andere einen Kennwert des Frequenzspektrums des Signals darstellen.
- 35
18. Vorrichtung nach einem der Ansprüche 6 bis 17, bei der die Eingabeeinrichtung dazu ausgelegt ist, für jede der Zeitperioden mehrere Bits zu erzeugen, die Bits aufweisen, die zwei oder mehr der folgenden Kennwerte darstellen:
- (i) Gesamtenergie innerhalb des Sprachfrequenzbands,
 - (ii) ein Maß für die Stimmhaftigkeit/Stimmlosigkeit,
 - (iii) die Energiepegel innerhalb einer Anzahl von jeweiligen Frequenzbändern,
 - (iv) der Unterschied zwischen den Pegeln in benachbarten, in (iii) genannten Bändern,
 - (v) die Unterschiede zwischen den Energiepegeln innerhalb einer Anzahl jeweiliger Frequenzbänder in aufeinanderfolgenden Abtastperioden,
 - (vi) die Anzahl der Nulldurchgänge der Sprachwellenform während einer gegebenen Periode,
 - (vii) die Anzahl der stationären Punkte in der Sprachwellenform während einer gegebenen Periode.
- 40
19. Vorrichtung nach Anspruch 18, bei der die Eingabeeinrichtung dazu ausgelegt ist, für jede Zeitperiode Bits zu erzeugen, die Werte entsprechend den Differenzen zwischen Energiepegeln in benachbarten, in (iii) genannten Bändern darstellen, wobei die Bänder logarithmisch beabstandete Mittenfrequenzen haben, und bei der für jede Zeitperiode jeder Wertesatz, betrachtet man sie in aufsteigender Reihenfolge hinsichtlich der Frequenz, vor Eintragung in das Bitmuster falls notwendig derart verschoben wird, daß der in der Reihenfolge am weitesten unten stehende Wert, der einen Energiepegel bei einem höheren Frequenzband darstellt, der gleich oder niedriger als der Energiepegel in einem niedrigeren Frequenzband darstellt, mit einem Datum versehen wird.
- 50
20. Vorrichtung nach einem der Ansprüche 6 bis 19, bei der zumindest ein typischer Wert des Sprachsignals in dem Muster durch zwei oder mehr Bits dargestellt wird, deren Bedeutung wechselweise voneinander abhängen, und wobei die Auswahl der eine Gruppe oder eine zusätzliche Gruppe bildenden Bits derart vorgenommen wird, daß diese zwei oder mehreren Bits für die gleiche Gruppe gewählt werden.
- 55



21. Vorrichtung nach einem der vorhergehenden Ansprüche, bei der das Ähnlichkeitsmaß die Bildung der Summe der abgerufenen Häufigkeiten umfaßt.
22. Vorrichtung nach einem der vorhergehenden Ansprüche, bei der das Ähnlichkeitsmaß die Bildung der Summe der Quadratwurzeln der abgerufenen Häufigkeiten umfaßt.
23. Vorrichtung nach einem der vorhergehenden Ansprüche, bei der jede Gruppe die gleiche Anzahl (n) von Bits hat.

Revendications

1. Dispositif de reconnaissance de formes comprenant :
 - (a) un moyen d'entrée pour délivrer une forme de bits à partir d'un signal d'entrée,
 - (b) un moyen de mémorisation d'affectation de données définissant des groupes de bits sélectionnés à partir de la forme,
 - (c) une mémoire de gabarit, et
 - (d) un moyen de commande pouvant être mis en oeuvre
 - (i) dans un mode apprentissage pour enregistrer dans la mémoire de gabarit, pour chacun des groupes définis, la fréquence à laquelle les combinaisons de bits à l'intérieur de ce groupe se produisent dans une pluralité desdites formes et pouvant être mis en oeuvre,
 - (ii) en mode reconnaissance en réponse à une forme inconnue délivrée par ledit moyen d'entrée, pour retrouver pour chaque groupe défini de bits sélectionnés à partir de ladite forme inconnue, en réponse aux bits du groupe, la fréquence s'il y en a une, qui est enregistrée dans la mémoire de gabarit pour cette combinaison particulière de bits, afin de former une mesure de la similitude avec cette forme inconnue en fonction des fréquences qui sont retrouvées et pour produire une sortie représentative de la reconnaissance satisfaisante, ou sinon en fonction du fait que la mesure de similitude satisfait un critère de reconnaissance, et
 - (iii) dans le mode reconnaissance, pour tout groupe quelconque pour lequel est déjà mémorisée dans la mémoire de gabarit, une fréquence pour une combinaison de bits particulière apparaissant dans la forme inconnue et pour lequel le critère de reconnaissance est satisfait, pour augmenter ladite fréquence mémorisée.
2. Dispositif selon la revendication 1 comportant une pluralité de mémoires de gabarit, le moyen de commande comportant un moyen de comparateur pour produire une sortie identifiant celle parmi les mémoires de gabarit pour laquelle la mesure de similitude satisfait le critère de reconnaissance.
3. Dispositif selon la revendication 1 ou la revendication 2, dans lequel, dans le mode reconnaissance, le moyen de commande peut être mis en oeuvre pour tout groupe quelconque, eu égard au fait qu'une fréquence n'est pas déjà mémorisée dans la mémoire de gabarit pour la combinaison de bits particuliers apparaissant dans la forme inconnue, s'il y en a une, suivant le fait que le critère de reconnaissance est satisfait, pour écrire une fréquence dans cette mémoire de gabarit se rapportant à ce groupe comportant cette combinaison.
4. Dispositif selon l'une quelconque des revendications précédentes, dans lequel le moyen de commande effectue l'ajustement de fréquences et l'entrée des fréquences s'il y en a, de manière conditionnelle, sur la présence ou l'absence d'un signal externe délivré au moyen de commande confirmant ou ne confirmant pas respectivement la reconnaissance satisfaisante.
5. Dispositif selon l'une quelconque des revendications précédentes, dans lequel la mémoire de gabarit ou chaque mémoire de gabarit a pour chaque groupe une section respective comprenant une pluralité d'entrées, chacune comprenant un emplacement pour la mémorisation d'une combinaison de bits particuliers et un emplacement pour la mémorisation de la fréquence associée à cette combinaison.
6. Dispositif selon l'une quelconque des revendications précédentes, dans lequel le moyen d'entrée peut être mis en oeuvre pour recevoir des signaux de parole et produire une forme de bits représentant, pour un temps de parole, des caractéristiques des signaux de parole dans des intervalles de temps successifs de cet intervalle.



7. Dispositif selon la revendication 6, dans lequel le dispositif est prévu pour rejeter les groupes obtenus exclusivement à partir des parties de la forme qui représentent un silence.
- 5 8. Dispositif selon la revendication 6 ou 7, dans lequel les données d'affectation représentent des groupes de bits sélectionnés de manière aléatoire à partir de la forme.
9. Dispositif selon la revendication 6 ou 7, dans lequel les données d'affectation représentent des groupes de bits, chacun desdits groupes comprenant des bits sélectionnés de la forme à l'intérieur d'une fenêtre temporelle limitée.
- 10 10. Dispositif selon la revendication 9, dans lequel la majorité des groupes sont chacun sélectionnés à partir de fenêtres temporelles différentes, de durée égale.
11. Dispositif selon la revendication 9 ou 10, dans lequel lesdits groupes sont divisés de façon imaginaire en une première catégorie et des catégories successives venant après chacune étant constituée de groupes obtenus à partir des fenêtres temporelles suivantes qui constituent les groupes de la catégorie précédente et le moyen de commande peut être mis en oeuvre au moins dans le mode reconnaissance, pour sélectionner eu égard à chaque groupe autre que ceux de la première catégorie, un ou plusieurs groupes supplémentaires, chacun étant constitué de bits décalés temporellement, d'une quantité qui est la même pour tous les bits de ce groupe par rapport à l'affectation des bits du groupe auquel il est apparenté, le nombre des groupes supplémentaires sélectionnés et la plage des décalages temporels auxquels ils sont soumis augmentant successivement pour les catégories successives.
- 15 12. Dispositif selon la revendication 8, dans lequel le moyen de commande peut être mis en oeuvre, au moins dans le mode reconnaissance pour sélectionner, en plus desdits groupes de bits, des groupes de bits supplémentaires chacun apparentés à l'un desdits groupes de bits et qui sont constitués de bits temporellement décalés par rapport aux affectations de bits du groupe apparenté, d'une quantité qui se rapporte à la position temporelle de ce bit à l'intérieur de la forme.
- 25 13. Dispositif selon la revendication 11 ou 12, dans lequel le moyen de commande peut être mis en oeuvre pour effectuer le processus supplémentaire de sélection de groupes de bits également dans le mode apprentissage après qu'au moins une forme ait été traitée dans le mode apprentissage.
- 30 14. Dispositif selon l'une quelconque des revendications 6 à 13, dans lequel le moyen de commande peut être mis en oeuvre, au moins dans le mode reconnaissance, afin d'effectuer un processus de positionnement temporel comprenant l'élaboration de la mesure de similitude de la forme avec une pluralité de positionnements temporels et le traitement de cette forme qui avec l'un quelconque des positionnements temporels produit la mesure de similitude qui indique la similitude la plus grande.
- 35 15. Dispositif selon la revendication 14, dans lequel le moyen de commande peut être mis en oeuvre pour effectuer le processus de positionnement temporel également dans le mode apprentissage après qu'au moins une forme ait été traitée dans le mode apprentissage.
- 40 16. Dispositif selon la revendication 14 ou 15, lorsqu'il dépend de la revendication 11, 12 ou 13, dans lequel le processus de positionnement temporel est limité à une partie antérieure de la forme.
- 45 17. Dispositif selon l'une quelconque des revendications 6 à 16, dans lequel le moyen d'entrée est prévu pour produire pour chacun desdits intervalles temporels une pluralité de bits dont certains représentent une caractéristique de la forme d'onde du signal de parole et dont les autres représentent une caractéristique du spectre de fréquence du signal.
- 50 18. Dispositif selon l'une quelconque des revendications 6 à 17, dans lequel le moyen d'entrée est prévu pour produire pour chacun desdits intervalles de temps une pluralité de bits comprenant des bits représentant deux ou plus des caractéristiques suivantes :
 - i) énergie totale à l'intérieur de la bande de fréquence de la parole
 - ii) mesure du son accentué/non accentué
 - 55 iii) les niveaux d'énergie à l'intérieur d'un certain nombre de bandes de fréquence respectives
 - iv) la différence entre les niveaux dans des bandes contiguës parmi les bandes référencées en (iii)
 - v) les différences entre les niveaux d'énergie à l'intérieur d'un certain nombre de bandes de fréquence



respectives dans des périodes d'échantillonnage successives

vi) le nombre de passages par zéro de la forme d'onde de la parole pendant une durée donnée

vii) le nombre de points stationnaires dans la forme d'onde de la parole pendant une durée donnée

- 5 **19.** Dispositif selon la revendication 18 dans lequel le moyen d'entrée est disposé pour produire pour chacun desdits intervalles de temps, des bits représentant les valeurs correspondant aux différences entre les niveaux d'énergie dans les bandes contiguës parmi les bandes référencées en (iii), lesdites bandes ayant des fréquences centrales espacées de façon logarithmique, et dans lequel pour chaque intervalle de temps, chaque ensemble de valeurs, considérées dans un ordre de fréquences ascendantes est, avant
- 10 l'entrée dans ladite forme de bit, décalée si nécessaire de sorte que la valeur d'ordre le plus faible qui représente le niveau d'énergie de bande de fréquence plus élevé est égal ou inférieur au niveau d'énergie de bande de fréquence plus faible, soit alignée avec une référence.
- 15 **20.** Dispositif selon l'une quelconque des revendications 6 à 19, dans lequel au moins une caractéristique du signal de parole est représentée dans la forme par deux bits ou plus, dont les significations sont interdépendantes et dans lequel la sélection des bits formant un groupe ou un groupe supplémentaire est telle que deux ou plus de ces bits soient sélectionnés pour le même groupe.
- 20 **21.** Dispositif selon l'une quelconque des revendications précédentes, dans lequel la mesure de similitude comporte l'élaboration de la somme des fréquences appariées.
- 25 **22.** Dispositif selon l'une quelconque des revendications précédentes, dans lequel la mesure de similitude comporte l'élaboration de la somme des racines carrées des fréquences appariées.
- 30 **23.** Dispositif selon l'une quelconque des revendications précédentes, dans lequel chaque groupe a le même nombre de bits (n).
- 35
- 40
- 45
- 50
- 55