



(72) SEJNOHA, VLADIMIR, US

(71) LERNOUT & HAUSPIE SPEECH PRODUCTS N.V., BE

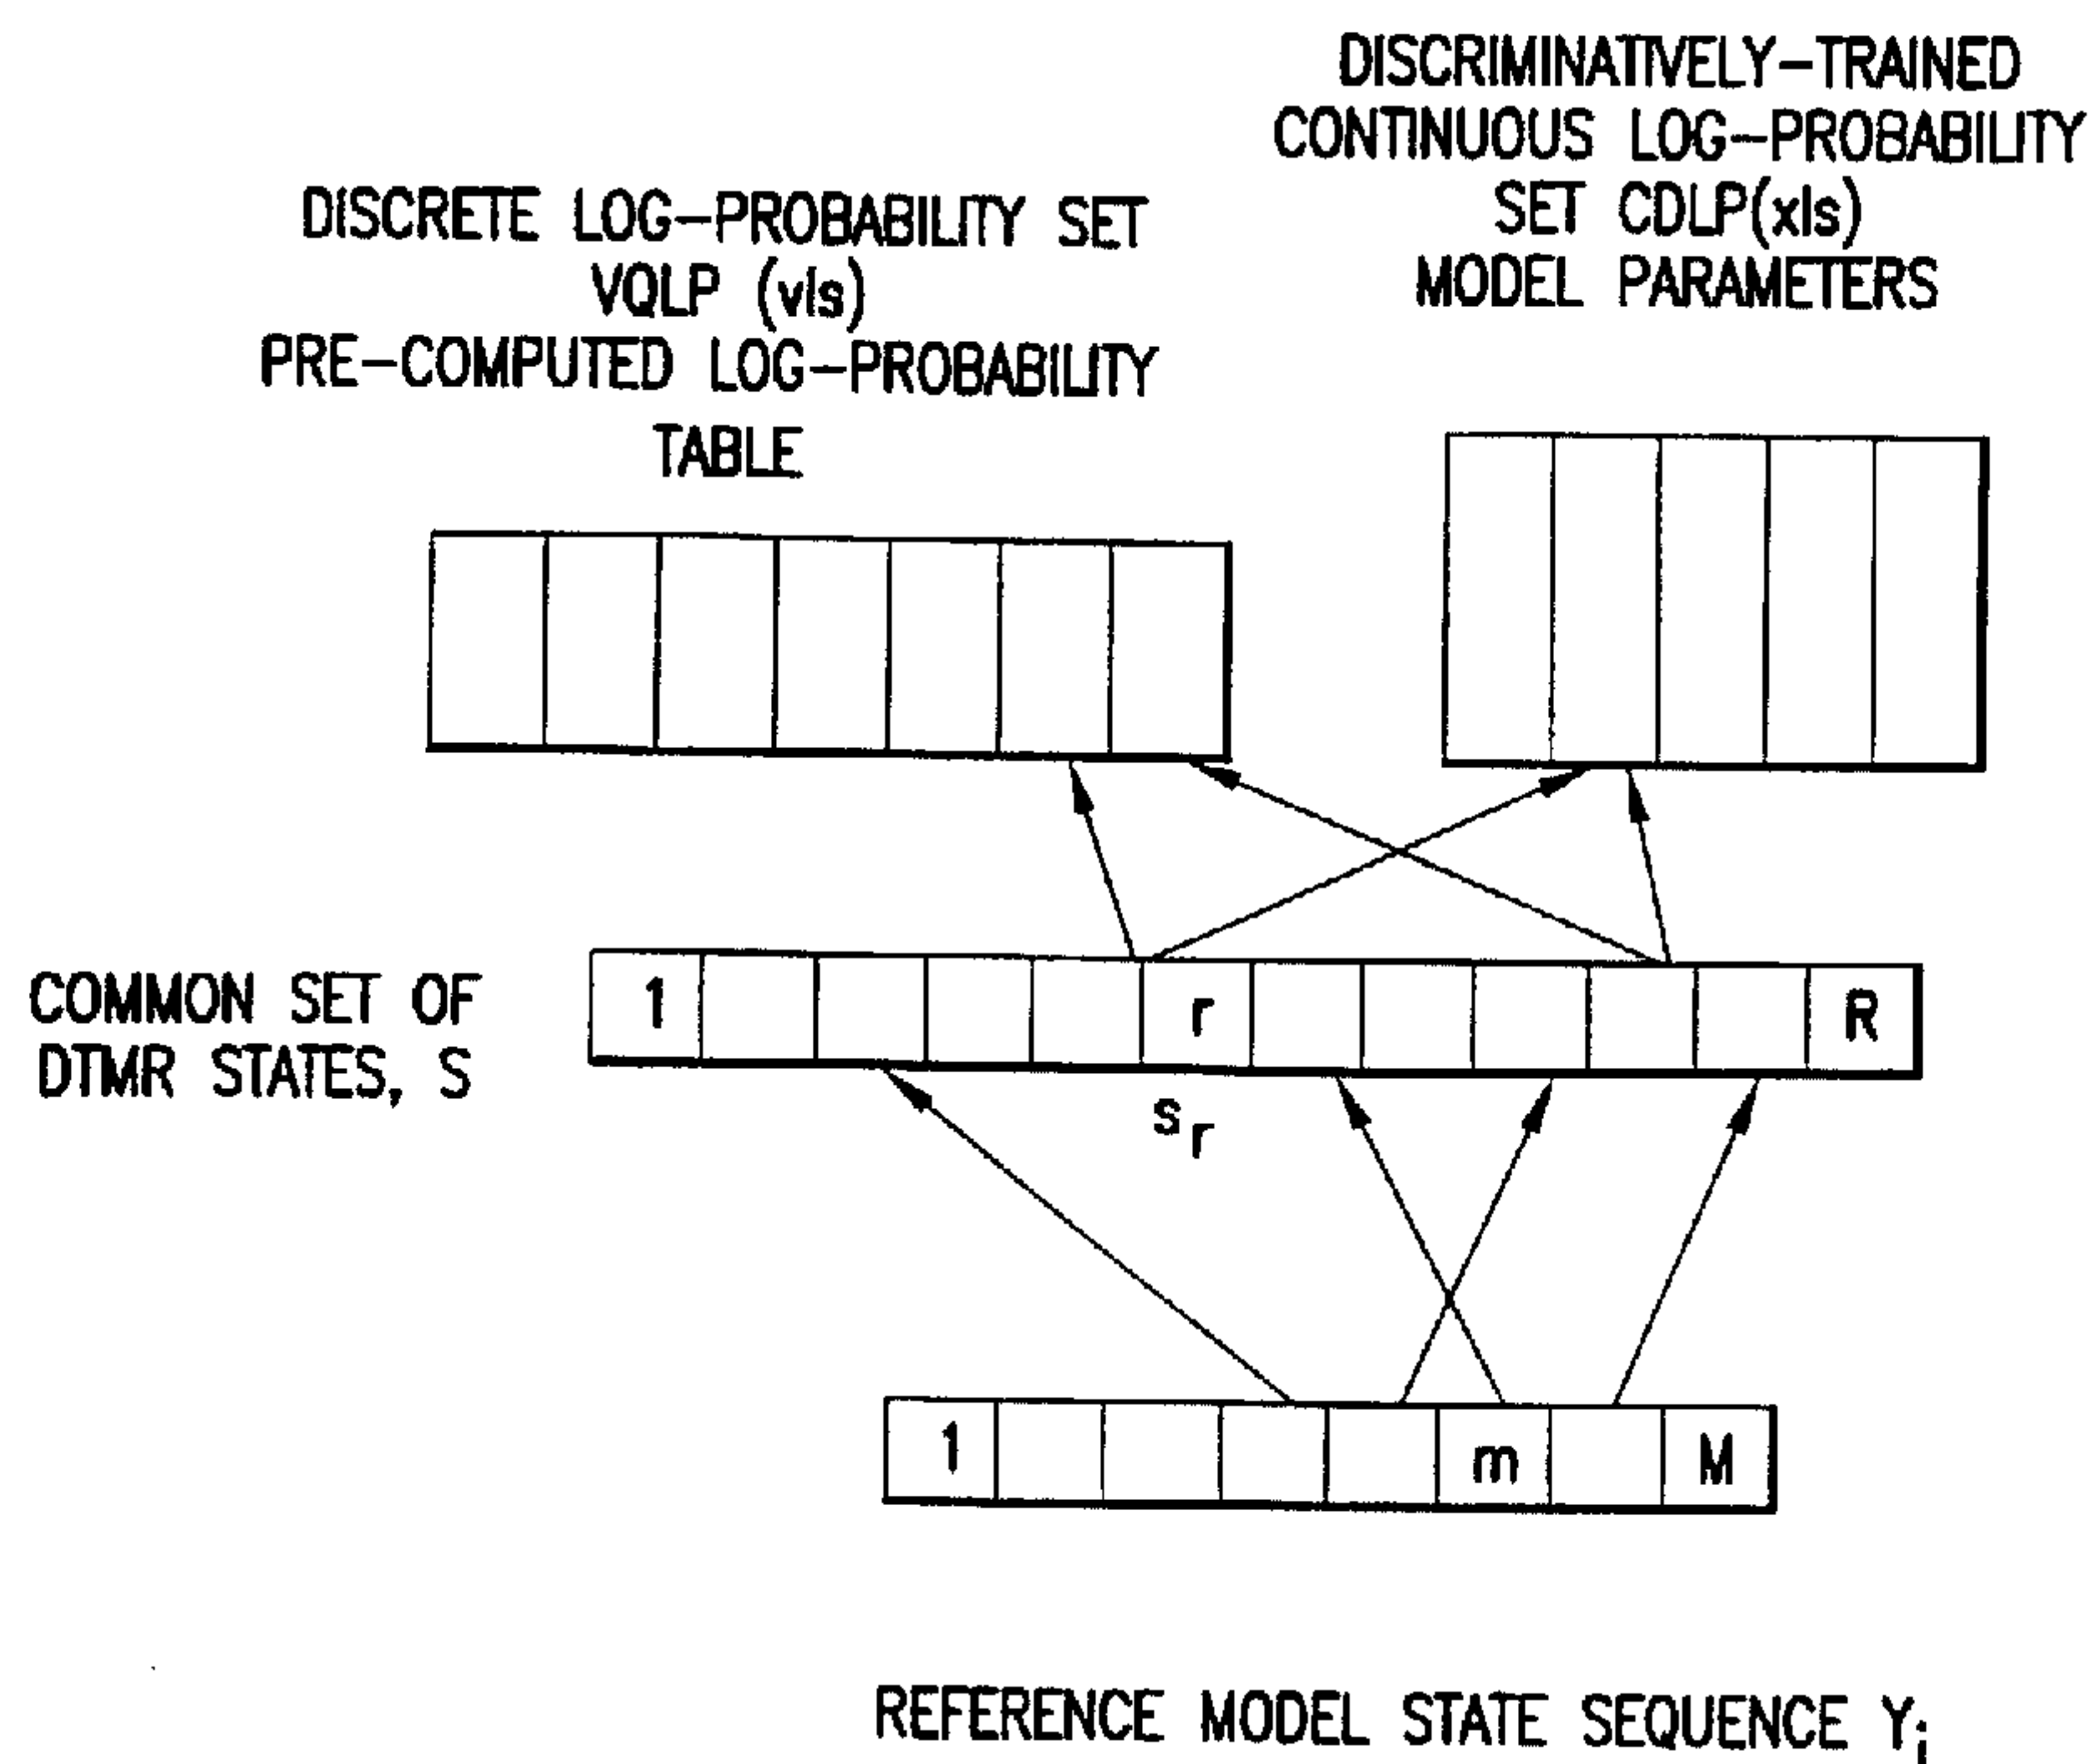
(71) SEJNOHA, VLADIMIR, US

(51) Int.Cl.<sup>6</sup> G10L 5/06

(30) 1997/03/14 (08/818,072) US

(54) **SYSTEME DE RECONNAISSANCE DE LA PAROLE FAISANT APPEL A DES MODELES ENTRAINES DE MANIERE DISCRIMINATIVE**

(54) **SPEECH RECOGNITION SYSTEM EMPLOYING DISCRIMINATIVELY TRAINED MODELS**



(57) L'invention concerne un système de reconnaissance de la parole dans lequel chaque émission d'entrée est convertie en une séquence de vecteurs bruts. Pour chaque vecteur brut, le système identifie, parmi plusieurs vecteurs quantifiés présélectionnés, celui qui correspond le mieux au vecteur brut. Les informations du vecteur brut sont cependant retenues pour être utilisées ultérieurement. Les modèles de mots de vocabulaire à

(57) In the speech recognition system disclosed herein, each input utterance is converted to a sequence of raw vectors. For each raw vector, the system identifies that one of a preselected plurality of quantized vectors which best matches the raw vector. The raw vector information is, however, retained for subsequent utilization. Each model of a vocabulary word to be recognized is in turn represented by sequence of states, the states being



reconnaître sont représentés, chacun à son tour par une séquence d'états, les états étant sélectionnés parmi un groupe d'états présélectionnés. Cependant, pour chaque état de modèle de mot, il existe une fonction de distribution de probabilité (pdf) discrète et une pdf continue, caractérisées par des paramètres présélectionnés réglables. Une table mise en mémoire contient des valeurs métriques de distance pour chaque combinaison de vecteur d'entrée quantifié avec un état de modèle caractérisé par les pdf discrètes. On aligne les modèles de mots avec une émission d'entrée en utilisant les pdf discrètes respectives et on génère les indices de correspondance initiaux au moyen de la table mise en mémoire. A partir de modèles de mots présentant une correspondance parfaite identifiés à partir des indices de correspondance initiaux, on génère un indice ordonné de ces modèles en utilisant les pdf continues respectives et les informations de vecteur brut. Après chaque émission, les paramètres de pdf continues présélectionnés sont ajustés de manière à augmenter, dans une proportion réduite, la différence d'indice entre le modèle du dessus et le modèle ordonné suivant. De préférence, si un utilisateur corrige un événement de reconnaissance antérieur en sélectionnant un modèle de mot différent à partir du groupe sélectionné correspondant, on réalise un réajustement des paramètres de pdf continues en effectuant, sur l'état de paramètres en cours, un ajustement opposé à celui effectué avec l'événement de reconnaissance original et en effectuant, sur l'état des paramètres en cours, un ajustement égal à celui qui aurait été effectué si le mot différent nouvellement identifié avait été le meilleur indice.

selected from a preselected group of states. However, for each word module state, there is provided both a discrete probability distribution function (pdf) and a continuous pdf characterized by preselected adjustable parameters. A stored table is provided which contains distance metric values for each combination of a quantized input vector with model state as characterized by the discrete pdfs. Word models are aligned with an input utterance using the respective discrete pdfs and initial match scores are generated using the stored table. From well matching word models identified from the initial match scores, a ranked scoring of those models is generated using the respective continuous pdfs and the raw vector information. After each utterance, the preselected continuous pdfs parameters are adjusted to increase, by a small proportion, the difference in scoring between the top and next ranking models. Preferably, if a user corrects a prior recognition event by selecting a different word model from the respective selected group, a re-adjustment of the continuous pdfs parameters is accomplished by performing, on the current state of the parameters, an adjustment opposite to that performed with the original recognition event and performing on the then current state of the parameters an adjustment equal to that which would have been performed if the newly identified different word model had been the best scoring.



CORRECTED  
VERSION\*CORRECTED  
VERSION\*\*

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

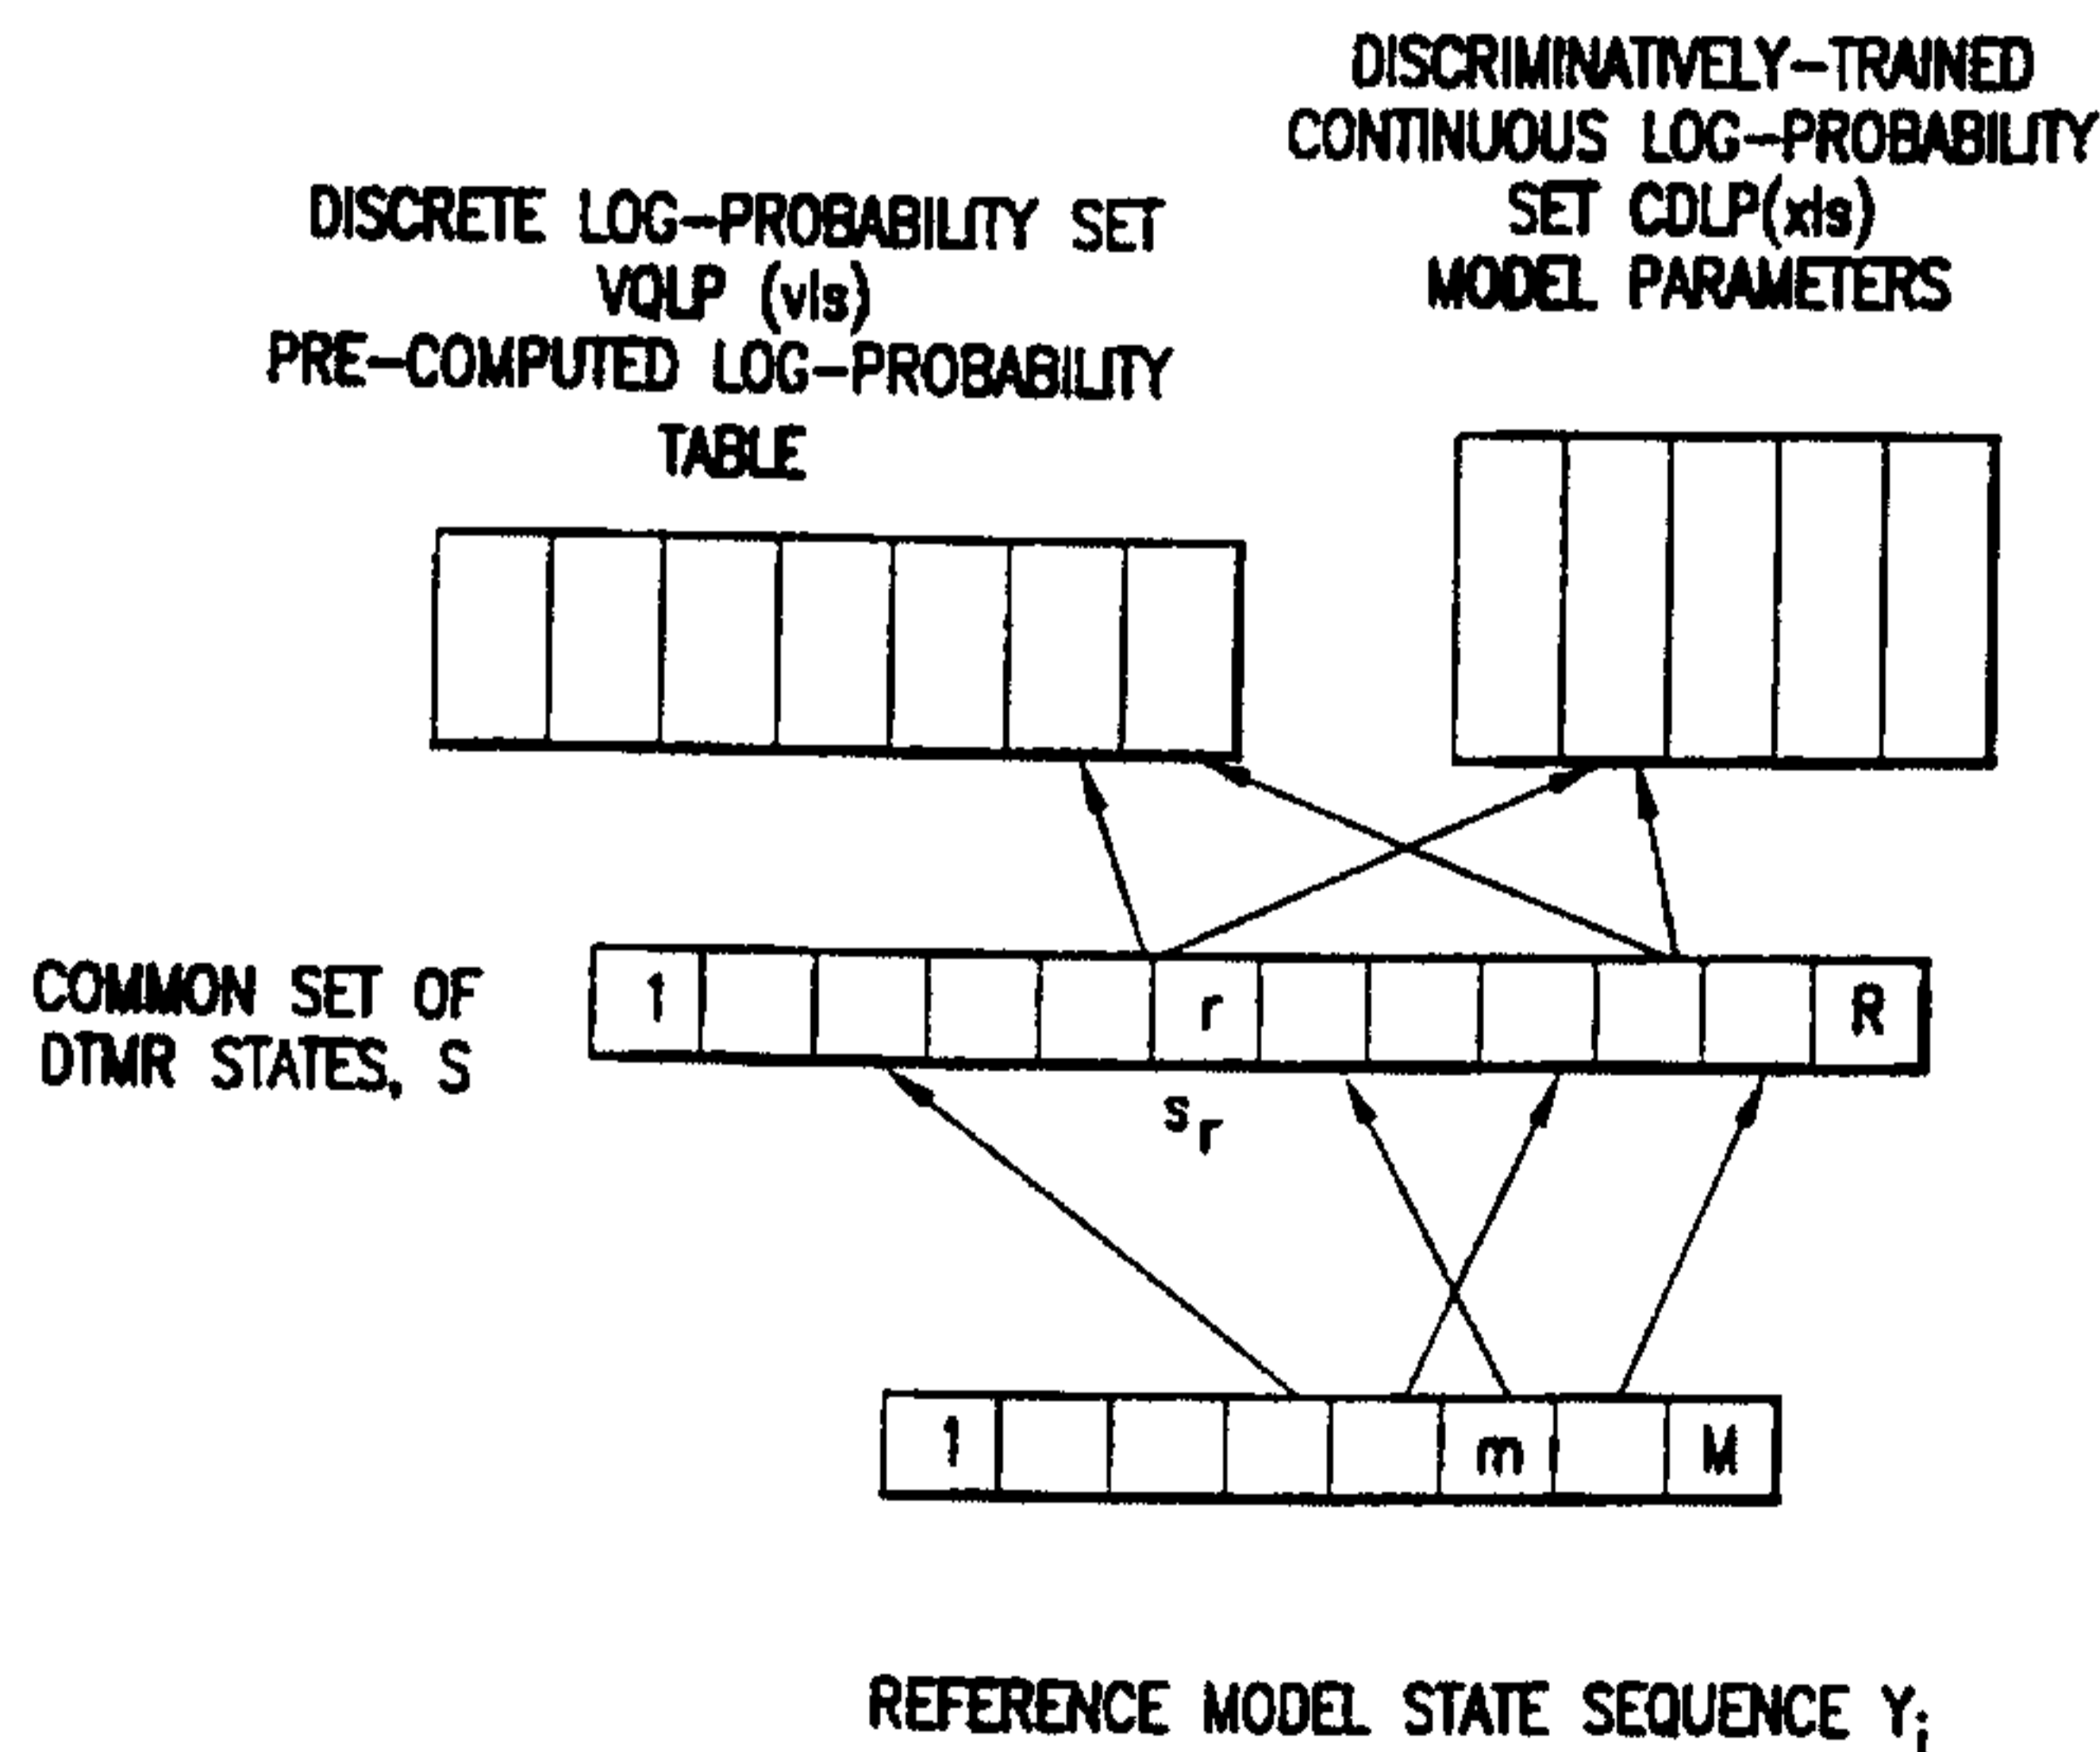
## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>6</sup> : <b>G10L 5/06</b>	<b>A1</b>	(11) International Publication Number: <b>WO 98/40876</b> (43) International Publication Date: 17 September 1998 (17.09.98)
(21) International Application Number: PCT/US98/04992 (22) International Filing Date: 13 March 1998 (13.03.98) (30) Priority Data: 08/818,072                      14 March 1997 (14.03.97)                      US (71) Applicant: LERNOUT & HAUSPIE SPEECH PRODUCTS N.V. [BE/BE]; Sint-Krispijnstaat 7, B-8900 Leper (BE). (71) (72) Applicant and Inventor: SEJNOHA, Vladimir [CA/US]; 29 Mt. Vernon Street #4, Cambridge, MA 02140 (US). (74) Agents: SUNSTEIN, Bruce, D. et al.; Bromberg & Sunstein LLP, 125 Summer Street, Boston, MA 02110-1618 (US).	(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>With international search report.</i>	

(54) Title: SPEECH RECOGNITION SYSTEM EMPLOYING DISCRIMINATIVELY TRAINED MODELS

## (57) Abstract

In the speech recognition system disclosed herein, each input utterance is converted to a sequence of raw vectors. For each raw vector, the system identifies that one of a preselected plurality of quantized vectors which best matches the raw vector. The raw vector information is, however, retained for subsequent utilization. Each model of a vocabulary word to be recognized is in turn represented by sequence of states, the states being selected from a preselected group of states. However, for each word module state, there is provided both a discrete probability distribution function (pdf) and a continuous pdf characterized by preselected adjustable parameters. A stored table is provided which contains distance metric values for each combination of a quantized input vector with model state as characterized by the discrete pdfs. Word models are aligned with an input utterance using the respective discrete pdfs and initial match scores are generated using the stored table. From well matching word models identified from the initial match scores, a ranked scoring of those models is generated using the respective continuous pdfs and the raw vector information. After each utterance, the preselected continuous pdfs parameters are adjusted to increase, by a small proportion, the difference in scoring between the top and next ranking models. Preferably, if a user corrects a prior recognition event by selecting a different word model from the respective selected group, a re-adjustment of the continuous pdfs parameters is accomplished by performing, on the current state of the parameters, an adjustment opposite to that performed with the original recognition event and performing on the then current state of the parameters an adjustment equal to that which would have been performed if the newly identified different word model had been the best scoring.



SPEECH RECOGNITION SYSTEM EMPLOYING DISCRIMINATIVELY  
TRAINED MODELS

5 Background of the Invention

The function of automatic speech recognition (ASR) systems is to determine the lexical identity of spoken utterances. The recognition process, also referred to as classification, typically begins with the conversion of an analog acoustical signal into a stream of digitally represented spectral vectors or frames which describe important characteristics of the signal at successive time intervals. The classification or recognition process is based upon the availability of reference models which describe aspects of the behavior of spectral frames corresponding to different words. A wide variety of models have been developed but they all share the property that they describe the temporal characteristics of spectra typical to particular words or subword segments. The sequence of spectral vectors arising from an input utterance is compared with the models and the success with which models of different words predict the behavior of the input frames, determines the putative identity of the utterance.

Currently most systems utilize some variant of a statistical model called the Hidden Markov Model (HMM). Such models consist of sequences of states connected by arcs, and a probability density function (pdf) associated with each state describes the likelihood of observing any given spectral vector at that state. A separate set of probabilities may be provided which determine transitions between states.

The process of computing the probability that an unknown input utterance corresponds to a given model, also known as decoding, is usually done in one of two standard ways. The first approach is known as Forward-Backward algorithm, and uses an efficient recursion to compute the match probability as the sum of the probabilities of

-2-

all possible alignments of the input sequence and the model states permitted by the model topology. An alternative, called the Viterbi algorithm, approximates the summed match probability by finding the single sequence of  $m$  states with the maximum probability. The Viterbi algorithm can be viewed as simultaneously performing an alignment of the input utterance and the model and computing the probability of that alignment.

HMMs can be created to model entire words, or alternatively, a variety of sub-word linguistic units, such as phonemes or syllables. Phone-level HMMs have the advantage that a relatively compact set of models can be used to build arbitrary new words given that their phonetic transcription is known. More sophisticated versions reflect the fact that contextual effects can cause large variations in the way different phones are realized. Such models are known as allophonic or context-dependent. A common approach is to initiate the search with relatively inexpensive context-independent models and re-evaluate a small number of promising candidates with context-dependent phonetic models.

As in the case of the phonetic models, various levels of modeling power are available in the case of the probability densities describing the observed spectra associated with the states of the HMM. There are two major approaches: the discrete pdf and the continuous pdf. In the former, the spectral vectors corresponding to the input speech are first quantized with a vector quantizer which assigns each input frame an index corresponding to the closest vector from a codebook of prototypes. Given this encoding of the input, the pdfs take on the form of vectors of probabilities, where each component represents the probability of observing a particular prototype vector given a particular HMM state. One of the advantages of this approach is that it makes no

-3-

assumptions about the nature of such pdfs, but this is offset by the information loss incurred in the quantization stage.

The use of continuous pdfs eliminates the quantization step, and the probability vectors are replaced by parametric functions which specify the probability of any  
5 arbitrary input spectral vector given a state. The most common class of functions used for this purpose is the mixture of Gaussians, where arbitrary pdfs are modeled by a weighted sum of Normal distributions. One drawback of using continuous pdfs is that, unlike in the case of the discrete pdf, the designer must make explicit assumptions about the nature of the pdf being modeled --- something which can be quite difficult since the  
10 true distribution form for the speech signal is not known. In addition, continuous pdf models are computationally far more expensive than discrete pdf models, since following vector quantization the computation of a discrete probability involves no more than a single table lookup.

The probability values in the discrete pdf case and the parameter values of the  
15 continuous pdf are most commonly trained using the Maximum Likelihood method. In this manner, the model parameters are adjusted so that the likelihood of observing the training data given the model is maximized. However, it is known that this approach does not necessarily lead to the best recognition performance and this realization has led to the development of new training criteria, known as discriminative, the objective of  
20 which is to adjust model parameters so as to minimize the number of recognition errors rather than fit the distributions to the data.

As used heretofore, discriminative training has been applied most successfully to small-vocabulary tasks. In addition, it presents a number of new problems, such as how to appropriately smooth the discriminatively-trained pdfs and how to adapt these

-4-

systems to a new user with a relatively small amount of training data.

To achieve high recognition accuracies, a recognition system should use high-resolution models which are computationally expensive (e.g., context-dependent, discriminatively-trained continuous density models). In order to achieve real-time  
5 recognition, a variety of speedup techniques are usually used.

In one typical approach, the vocabulary search is performed in multiple stages or passes, where each successive pass makes use of increasingly detailed and expensive models, applied to increasingly small lists of candidate models. For example, context  
independent, discrete models can be used first, followed by context-dependent  
10 continuous density models. When multiple sets of models are used sequentially during the search, a separate simultaneous alignment and pdf evaluation is essentially carried out for each set.

In other prior art approaches, computational speedups are applied to the evaluation of the high-resolution pdfs. For example, Gaussian-mixture models are  
15 evaluated by a fast but approximate identification of those mixture components which are most likely to make a significant contribution to the probability and a subsequent evaluation of those components in full. Another approach speeds up the evaluation of Gaussian-mixture models by exploiting a geometric approximation of the computation. However, even with speedups the evaluation can be slow enough that only a small  
20 number can be carried out.

In another scheme, approximate models are first used to compute the state probabilities given the input speech. All state probabilities which exceed some threshold are then recomputed using the detailed model, the rest are retained as they are. Given the new, composite set of probabilities a new Viterbi search is performed to determine the

-5-

optimal alignment and overall probability. In this method, the alignment has to be repeated, and in addition, the approximate and detailed probabilities must be similar, compatible quantities. If the detailed model generates probabilities which are significantly higher than those from the approximate models the combination of the two will most likely not lead to satisfactory performance. This requirement constrains this method to use approximate and detailed models which are fairly closely related and thus generate probabilities of comparable magnitude. It should also be noted that in this method there is no guarantee that all of the individual state probabilities that make up the final alignment probability come from detailed models.

10           The present invention represents a novel approach to the efficient use of high-resolution models in large vocabulary recognition. The proposed method benefits from the use of a continuous density model and a discriminative training criterion which leads to a high recognition performance on a large vocabulary task at the cost of only a marginal increase of computation over a simple discrete pdf system. Another novel  
15 feature of the new approach is its ability to make use of limited quantities of new data for rapid adaptation to a particular speaker.

As was mentioned above, the probability that an input utterance corresponds to a given HMM can be computed by the Viterbi algorithm, which finds the sequence of model states which maximizes this probability. This optimization can be viewed as a  
20 simultaneous probability computation and alignment of the input utterance and the model.

In accordance with one aspect of the present invention, it has been determined that the alignment paths obtained with relatively computationally inexpensive discrete pdf models can

-6-

be of comparable quality to those obtained with computationally costly continuous density pdf models, even though the match probabilities or metrics generated by the discrete pdf alignment do not lead to sufficiently high accuracy for large vocabulary recognition.

5 In accordance with another aspect of the invention, there is provided a decoupling of the alignment and final probability computation tasks. A discrete-pdf system is used to establish alignment paths of an input utterance and a reference model, while the final probability metric is obtained by post-processing frame-state pairs with more powerful, discriminatively trained continuous-density pdfs, but using the same  
10 alignment path.

Unlike conventional systems, where model states are characterized by one particular type of observed pdf, the state models in the present system are thus associated with both a discrete (low-resolution) pdf and a discriminatively trained, continuous-density (high-resolution) pdf. The high-resolution pdfs are trained using  
15 alignments of models and speech data obtained using the low-resolution pdfs, and thus the discriminative training incorporates knowledge of the characteristics of the discrete pdf system.

#### Brief Description of the Invention

In the speech recognition system of the present invention, each input  
20 utterance is converted to a sequence of raw or unquantized vectors. For each raw vector the system identifies that one of a preselected plurality of quantized vectors which best matches the raw vector. The raw vector information is, however, retained for subsequent utilization. Each word model is represented by a sequence of states, the states being selected from a preselected group of states. However, for

-7-

each word model state, there is provided both a discrete probability distribution function (pdf) and a continuous pdf characterized by preselected adjustable parameters. A stored table is provided which contains distance metric values for each combination of a quantized input vector with model state as characterized by  
5 the discrete pdfs.

Word models are aligned with an input utterance using the respective discrete PDFs and initial match scores are generated using the stored table. From well matching word models identified from the initial match scores, a ranked scoring of those models is generated using the respective continuous pdfs and the  
10 raw vector information. After each utterance, the preselected parameters are adjusted to increase, by a small proportion, the difference in scoring between the top and next ranking models.

Preferably, if a user corrects a prior recognition event by selecting a different word model from the respective selected group, a re-adjustment of the  
15 continuous pdf parameters is accomplished by performing, on the current state of the parameters, an adjustment opposite to that performed with the original recognition event and performing on the then current state of the parameters an adjustment equal to that which would have been performed if the newly identified different word model had been the best scoring.

20

#### Brief Description of the Drawings

Fig. 1 is a block diagram of a speech recognition system in accordance with the present invention;

Fig. 2 illustrates vocabulary word models used in the speech recognition system of the present invention;

-8-

Fig. 3 illustrates a recursion procedure used in the speech recognition system of the present invention;

Fig. 4 illustrates a training data structure set used in training word models;

Fig. 5 is a flow chart illustrating initial, batch training of word models; and

5 Fig. 6 is a flow chart illustrating on-line adaptive training of word models.

Corresponding reference characters indicate corresponding elements throughout the several views of the drawings.

#### Description of the Preferred Embodiment

As indicated previously, the present invention is particularly concerned with  
10 the provision of discriminatively trained multi-resolution vocabulary models which increase accuracy and reduce computational load in an automatic speech recognition (ASR) system. At the outset, however, it is appropriate to describe in general terms the type of speech recognition system to which the present invention is applicable.

15 Referring now to Figure 1, the computer system illustrated there is of the type generally referred to as a personal computer. The computer runs under the MS DOS or WINDOWS® operating system and is organized around a system bus, designated generally by reference character 11. The system bus may be of the so called EISA type (Extended Industry Standards Association). The computer system  
20 utilizes a microprocessor, designated by reference character 13, which may, for example, be an Intel Pentium type processor. The system is also provided with an appropriate amount of local or random access memory, e.g., 32 megabytes, designated by reference character 15. Additional storage capacity is provided by a hard disk 17 and floppy diskette drive 19 which operate in conjunction with a

-9-

controller 23 which couples them to the system bus.

User input to the computer system is conventionally provided by means of keyboard 25 and feedback to the user is provided by means of a CRT or other video display 27 operating from the bus through a video controller 29. External  
5 communications may be provided through an I/O system designated by reference character 31 which supports a serial port 33 and a printer 35. Advantageously, a fax modem may be provided as indicated by reference character 37. This is particularly useful for forwarding structured medical reports as described in co-assigned U. S .  
Patent No. 5,168,548.

10 To facilitate the use of the computer system for speech recognition, a digital signal processor is provided as indicated by reference character 16, typically this processor being configured as an add-in circuit card coupled to the system bus 11. As is understood by those skilled in the art, the digital signal processor takes in analog signals from a microphone, designated by reference character 18, converts  
15 those signals to digital form and processes them e.g., by performing a Fast Fourier Transform (FFT), to obtain a series of spectral frames or vectors which digitally characterize the speech input at successive points in time. As used herein, these input vectors are referred to as the raw input vectors. In the embodiment being described, acoustic vectors ( $X_n$ ) are generated at a rate of one every 10 ms, and have  
20 14 output dimensions.

Preferably, the raw vectors are subjected to a gender-normalizing linear discriminant analysis, as described in my co-pending, coassigned application Serial No. 08/185,500, the disclosure of which is incorporated herein by reference. The purpose of the analysis is to transform the spectral frames so as to enhance the

-10-

discriminability of different phonetic events. While the raw vectors are subsequently quantized for use in alignment and initial scoring, the data comprising the raw vectors is preserved for use in more precise final scoring using continuous pdfs as described hereinafter.

5 Thus  $X_u = (x_{u,1}, \dots, x_{u,t}, \dots, x_{u,T_u})$ , where  $T_u$  is the length and  $X_{u,t}$  is the  $t^{\text{th}}$  vector of size 14 in the  $u^{\text{th}}$  input utterance.

The transformed acoustic frames are vector quantized with a codebook of 1024 standard vector prototypes and each original spectral frame  $x_t$  (omitting the subscript  $u$ ) is assigned a corresponding vector quantizer (VQ) label  $v_r$ . Each  
10 sequence  $\mathbf{X}$  thus gives rise to a VQ label sequence  $\mathbf{V} = (v_1, \dots, v_t, \dots, v_T)$ .

Reference vocabulary models are composed of sequences of states

$\mathbf{Y}_i = (y_{i,1}, \dots, y_{i,m}, \dots, y_{i,M_i})$ , where  $M_i$  is the length of a model and  $i$  is the model index.

Each model state  $y_{i,m}$  is a pointer into a common set of  $R$  DTMR states,  $\mathbf{S} = (s_1, \dots, s_r, \dots, s_R)$ , each of which is associated in turn with two distinct types of pdf  
15 selected from two common pdf pools.

The first type of pdf pool contains discrete distributions which express the probability of observing a quantized frame  $v_t$ , given a state  $s_r$  referenced by  $y_{i,m}$  which occurs at the  $m^{\text{th}}$  position in the  $i^{\text{th}}$  model, *i.e.*,  $\Pr(v_t | y_{i,m}) = \Pr(v_t | s_r)$ . The computation of the match probability is simplified if the pdfs are converted to  
20 negative logarithms and thus we define the quantity  $\text{VQLP}((v_t | y_{i,m})) = -\log(\Pr(v_t | y_{i,m}))$ . Note that VQLP is essentially a table of precomputed log-probabilities and thus the evaluation of the discrete-pdf models consists of a very fast table lookup.

The second pool of pdfs, on the other hand, is made up of continuous distributions which give the probability of observing a specific spectrum  $\mathbf{X}_t$  given a

-11-

particular state  $s_r$  referenced by  $y_{i,m}$ , *i.e.*,  $\Pr(\mathbf{X}_t|y_{i,m}) = \Pr(\mathbf{X}_t|s_r)$ . As is the case for the discrete pdfs it is more convenient to use the continuous-density probabilities in the log-domain, and thus we define  $\text{CDLP}(\mathbf{X}_t|y_{i,m}) = -\log(\Pr(\mathbf{X}_t|y_{i,m}))$ .

The continuous pdfs are parametric models and thus the probabilities cannot be precomputed. Rather than storing pre-computed probabilities as is the case for the discrete pdfs, we store the pdf parameters themselves and use them to compute the log-probabilities for specific input frames. Note that individual palls in each set may be shared by acoustically similar states in different models. Fig. 2 illustrates the relationship between model states and the two sets of log-pdfs.

The vector-quantized input utterances  $\mathbf{V}$  are matched against reference models  $\mathbf{Y}_i$  by the Viterbi algorithm described in greater detail hereinafter using the discrete-pdf section of the DTMR models. The algorithm evaluates  $-\log(\Pr(\mathbf{V}|\mathbf{Y}_i))$ , the negative logarithm of the probability of observing the input sequence given the model by finding the best alignment path between the input utterance and the model states.

Define  $\text{Sum}(t,m)$  as the accumulated negative log-probabilities.

The alignment algorithm used in this work can then be summarized as follows.

```

Initialization:
Sum(0, 0) = 0
20
Evaluation:
for t := 1 to T do
  for m := 1 to  $M_i$  do
    Sum(t, m) =  $\text{VQLP}(v_t, y_{i,m}) + \min(\text{Sum}(t-1, m), (\text{Sum}(t-1, m-1), (\text{Sum}(t, m-1)))$ 
25    Predt(t, m) = arg mint(Sum(t-1, m), (Sum(t-1, m-1), (Sum(t, m-1))
    Predm(t, m) = erg minm(Sum(t-1, m), (Sum(t-1, m-1), (Sum(t, m-1))
  end;
end;

30 Termination:
return Sumi = Sum(T,M)/(T+M)

```

-12-

where  $\text{Pred}_i(t, m)$  and  $\text{Pred}_m(t, m)$  are the indices of the best predecessor score at position  $t, m$ .

This basic recursion is also illustrated in Fig. 3 of the drawings.

5           The following structures are set up to store the alignment path between an input utterance and a given reference model:

$$\mathbf{f}_i = (f_{i,1}, \dots, f_{i,p}, \dots, f_{i,P_i}) \quad (1)$$

$$\mathbf{q}_i = (q_{i,1}, \dots, q_{i,p}, \dots, q_{i,P_i}) \quad (2)$$

where  $f_{i,p}$  is the input frame index and  $q_{i,p}$  is the state index at position  $p$  on the path

10          for the  $i^{\text{th}}$  reference model, and  $P_i$  is the path length. The best alignment path is recovered by using the predecessor arrays  $\text{Pred}_i(t, m)$  and  $\text{Pred}_m(t, m)$  in the following backtracking recursion:

Initialization:

$p = 0, t = T, m = M_i$

15

Backtracking:

while  $t \geq 1$  and  $m \geq 1$  do

$f_{i,p} = t$

$q_{i,p} = y_{i,m}$

20

$t_{\text{new}} = \text{Pred}_i(t, m)$

$m_{\text{new}} = \text{Pred}_m(t, m)$

$t = t_{\text{new}}$

$m = m_{\text{new}}$

$p++$

25

end;

Termination:

$P_i = p$

30

The original acoustic vector at a particular path point  $p$  can thus be identified as

$\mathbf{x}(f_{i,p})$  while the state index at path position  $p$  is directly given by  $q_{i,p}$ . The

normalized scores  $\text{Sum}_i$  for all the reference models aligned with a particular input

utterance can be sorted and only a small number of models with the lowest

35          normalized score need be restored in the next recognition stage.

-13-

In the rescoring stage the alignment paths for a small number of the models with the best discrete-pdf scores are traversed, and new scores are computed for the frame-state pairs defined by these paths using the set of discriminatively trained continuous density pdfs.

5 The continuous density pdfs used in this work are a simplified form of Gaussian Mixtures. Experimental evidence revealed that with the use of discriminative training there was no advantage to using the full mixture models over the simplified version. In addition, reducing the number of free parameters in the model significantly improves their trainability with limited quantities of data.

10 The standard Gaussian Mixture log-probability density function GMLP is defined as follows:

$$GMLP(x(t), s_r) = -\log\left(\sum_k^{N(s_r)} a(s_r, k) N(x(t); \mu(s_r, k); \Sigma(s_r, k))\right) \quad (3)$$

where  $a(s_r, k)$  is the weight of mixture component  $k$  in state  $s_r$ , and  $N(x; \mu; \Sigma)$  denotes the probability of observing  $x(t)$  given a multivariate Gaussian with mean  $\mu$  and covariance  $\Sigma$ .  $N(s_r)$  is the number of mixture components.

The discriminatively trained continuous density log-pdf (CDLP) used in this work is as follows

$$CDLP(x(f_{i,p}), q_{i,p}) = \min_{1 \leq k \leq N(q_{i,p})} [d(x(f_{i,p}), \mu(q_{i,p}, k))] \quad (4)$$

where

$$20 \quad d(x(f_{i,p}), \mu(q_{i,p}, k)) = \sum_{l=1}^{14} (x(f_{i,p}, l) - \mu(q_{i,p}, k, l))^2 \quad (5)$$

The continuous pdf model for state  $q_{i,p}$  thus consists of  $N(q_{i,p})$  14-dimensional mean vectors  $\mu$ . Due to the lack of normalizing terms in equation (4), CDLP is not a true log-probability, and thus is not interchangeable with the

-14-

discrete log-probabilities VQLP. This incompatibility is not an issue, however, because once the alignment paths are established the discrete log-probabilities are no longer used. The ability to utilize incompatible pdfs constitutes an advantage over known schemes.

5           The new score for a path corresponding to an alignment of input utterance with reference model  $i$  is obtained as

$$D_i = \frac{1}{P_i} \sum_{p=1}^{P_i} CDLP(x(f_{i,p}, q_{i,p})) \quad (6)$$

The rescored models are then re-sorted according to their new scores.

10           The role of the discrete-density component of the DTMR models is two-fold: for each input utterance it is used to screen out the great majority of incorrect models and produce a small set of likely candidates, and it is also used to obtain accurate alignment of the input utterance to reference models. It is, however, not called upon to provide fine discrimination between highly confusable models. Rather, that is the role of the continuous density rescoreing pdfs.

15           For these reasons it is sufficient to rely on conventional Maximum-Likelihood training for the discrete-density component, and apply a discriminative criterion to the training of the continuous density component only. The continuous pdf training however, uses alignment paths established on the basis of the discrete pdfs.

20           The first step in the training of the continuous density pdfs is the initialization of the mean vectors  $\mu_{s,k}$ . This can be done by training a conventional Maximum Likelihood Gaussian Mixture pdf for each model state from the input utterance frames aligned with that state using the discrete-pdf

-15-

component. The total number of mean vectors can be set to reflect the variance of the data frames aligned with each state during the iterative training. Upon convergence of the initial training, the mean vectors with significant probabilities are retained, while all other parameters associated with a standard Gaussian Mixture model are discarded.

The next step consists of the discriminative training of the mean vectors. This is accomplished by defining an appropriate training objective function which reflects recognition error-rate and optimizing the mean parameters so as to minimize this function.

One common technique applicable to the minimization of the objective function is gradient descent optimization. In this approach, the objective function is differentiated with respect to the model parameters, and the parameters are then modified by the addition of the scaled gradient. A new gradient which reflects the modified parameters is computed and the parameters are adjusted further. The iteration is continued until convergence is attained, usually determined by monitoring the performance on evaluation data independent from the training data.

A training database is preprocessed by obtaining for each training utterance a short list of candidate recognition models. Each candidate list contains some number of correct models (subset C) and a number of incorrect (subset I) models. Each list is sorted by the score  $D_i$ , and an augmented alignment path structure is retained for each reference model in the list. The additional stored path information is as follows:

-16-

$$\mathbf{b}_i = (b_{i,1}, \dots, b_{i,p}, \dots, b_{i,P_i}) \quad (7)$$

$\mathbf{b}_i$  is used to store the index of the best mean vector at a particular path point. For example, if  $p$  connects the frame  $\mathbf{x}(f_{i,p})$  and state  $\mathbf{q}_{i,p}$ ,

$$b_{i,p} = \arg \min_{1 \leq k \leq N(q_{i,p})} [d(x(f_{i,p}), \mu(q_{i,p}, k))] \quad (8)$$

5 Fig. 4 illustrates the training structure set for an input utterance. An error function  $\epsilon_u$  for a particular training utterance  $u$  is computed from the pairwise error functions  $o_{i,j}$ :

$$\epsilon_u = \frac{1}{2} \sum_{i \in C} \sum_{j \in I} o_{i,j}^2 \quad (9)$$

where

$$10 \quad o_{i,j} = (1 + e^{-\beta(D_i - D_j)})^{-1} \quad (10)$$

$\beta$  is a scalar multiplier,  $D_i$ ,  $i \in C$  is the alignment score of the input token and a correct model  $i$ , and  $D_j$ ,  $j \in I$  is the score of the token and an incorrect model  $j$ . The sizes of the sets  $C$  and  $I$  can be controlled to determine how many correct models and incorrect or potential intruder models are used in the training.

15  $o_{i,j}$  takes on values near 1 when the correct model score  $D_i$  is much greater (i.e., worse) than the intruder score  $D_j$ , and near 0 when the converse is true. Values greater than 0.5 represent recognition errors while values less than 0.5 represent correct recognitions. The parameter  $\beta$  controls the amount of influence "near-errors" will have on the training.

20 The score  $D_i$  between the training utterance and the target model  $i$  is obtained by rescoreing the alignment path as shown in equation (6). In practice the normalization by path length  $P_i$  can be ignored during training. Thus

$$D_i = \sum_{p=1}^{P_i} CDLP(x(f_{i,p}), q_{i,p}) \quad (11)$$

-17-

which can be rewritten as:

$$D_i = \sum_{p=1}^{P_i} d(x(f_{i,p}), \mu(q_{i,p}, b_{i,p})) \quad (12)$$

A similar expression can be written for  $D_j$ .

Differentiating the error function with respect to a particular component of

5 the mean vector  $\mu(s, k, l)$  yields:

$$-\frac{\partial \epsilon_u}{\partial \mu(s, k, l)} = 2\beta \sum_{i \in C} \sum_{j \in I} o_{ij}^2 (1 - o_{ij}) \left\{ \sum_p^{P_i} (x(f_{i,p}, l) - \mu(q_{i,p}, b_{i,p}, l)) \delta(s, q_{i,p}) \delta(k, b_{i,p}) - \sum_p^{P_j} (x(f_{j,p}, l) - \mu(q_{j,p}, b_{j,p}, l)) \delta(s, q_{j,p}) \delta(k, b_{j,p}) \right\} \quad (13)$$

where  $\delta(a, b)$  is the Kronecker delta and equals 1 if  $a = b$ , and 0 otherwise. The

gradient is averaged over all utterances and correct-incorrect pairs:

$$10 \quad \Delta \mu(s, k, l) = \frac{1}{U} \sum_u \frac{1}{N_{C,I,u}} \frac{-\partial \epsilon_u}{\partial \mu(s, k, l)} \quad (14)$$

where  $N_{C,I,u}$  is the number of correct-incorrect model pairs for utterance  $u$ . The

mean components are modified by the addition of the scaled gradient:

$$\hat{u}(s, k, l) = \mu(s, k, l) + \omega \Delta \mu(s, k, l) \quad (15)$$

where  $w$  is a weight which determines the magnitude of the change to the

15 parameter set in one iteration. This procedure is illustrated in the flowchart of Fig.

5.

Initially, candidate models are selected using the discrete density pdfs as indicated in step 101. Again using the discrete pdfs, the input utterances aligned with the best models using the Viterbi algorithm and the traceback information is stored as indicated at block 103. The scores are sorted as indicated at block 105 and then the top scoring models are re-scored as indicated at block 107 using the continuous density pdfs, the rescoring being done along the alignment path determined with the discrete pdfs.

20

-18-

As indicated at block 109, the models are then re-sorted based on the scores obtained with the continuous density pdfs.

Correct and incorrect models are identified as indicated at block 111 and for each pair of correct and incorrect models an error function is computed as indicated at block 113. Since multiple models may be used for each vocabulary word, the procedure provides for subsets rather than just individual correct and incorrect examples.

As indicated at block 115, a gradient is accumulated for each pair trace backed along the correct and incorrect paths. An accumulated gradient is applied to the continuous density pdf parameters as indicated at block 119. A test for convergence is applied as indicated at block 121 and the procedure beginning at block 117 is repeated until the models have converged.

In each iteration of the *batch mode* training of the DTMR models outlined in the previous section, all training utterances are processed before the model parameters are updated. It is however also possible to train the models with an on-line adaptive algorithm, where the models are updated after each training utterance has been processed. The on-line training makes it possible to rapidly adapt the DTMR models with limited amounts of speech from a new user of the recognition system, and ideally this may be done in a fashion invisible to the user, with speech produced in the course of doing useful work with the recognition system.

Like batch training, on-line training requires the computation of the error function gradient for all current model parameters specified by the correct and

-19-

incorrect alignment paths in the candidate set for a particular input utterance  $u$ .

Unlike in the batch training case, the gradient is not accumulated but is applied immediately to the model parameters:

$$\hat{u}(s, k, l) = \mu(s, k, l)_{u-1} + \omega' \Delta_u \mu(s, k, l)_{u-1} \quad (16)$$

- 5 The notation  $\Delta u$  means that the utterance  $u$  is used to compute the gradient, and the operation is performed on the current model  $\mu(s, k, l)_{u-1}$  (which was presumably also adapted on the previous utterance  $u - 1$ ).

The weighting used in the on-line adaptation  $\omega'$  is set much smaller than the weighting used in the batch-mode training since the reliability of the change  
10 estimated from a single utterance is considerably lower than the estimate from a complete training set. In other words,  $\omega' \ll \omega$ .

A complicating factor in on-line adaptation is that the identity of the input utterances is not known with certainty. Relying on the recognition system to  
15 identify the input utterances will inevitably lead to errors and misadaptations of the models. Delaying the adaptation to give the user a chance to make corrections is not desirable, because given the different correction strategies favored by different users, it is difficult to predict how long the delay needs to be.

The solution to this problem provided by the present invention is to begin  
20 by assuming that the top-choice recognition candidate is in fact the correct answer and to update the models immediately. However, if the user makes a correction at some subsequent time, the original misadaptation will be undone and a new modification of the model parameters will be performed based on the corrected information.

-20-

In order to undo the incorrect adaptation at a later time, the original input utterance corresponding to each candidate set must be retained, although the candidate alignment paths need not be saved. When the user specifies the correct answer, the candidate alignment paths are regenerated and the utterance gradient term is recomputed. The weighted gradient is subtracted from the affected model parameters. A new gradient term, reflecting the correct target model is calculated and applied to the DTMR parameters. This sequence of operations does not completely undo the original error because to do so would mean undoing all the other updates that may have intervened between the original recognition and the correction action. In practice, however, the delayed correction has proved to be as effective as supervised adaptation, i.e. where the correct answer is identified before gradient computation.

The delayed correction algorithm is as follows:

```

15  for  $u := 1$  to  $U$  do

    Obtain candidate list for  $u$  using  $\mu_{u-1}$ 
    Identify subsets  $I_{top-choice}$  (assumed incorrect models)
    and  $C_{top-choice}$  (assumed correct models).
20  Compute  $\Delta_u \mu(s, k, l)_{u-1}$  for all  $s, k, l$  specified by alignment
    paths for all pairs in  $I_{top-choice}$  and  $C_{top-choice}$ 
    Update  $\mu(s, k, l)_u = \mu(s, k, l)_{u-1} + \omega' \Delta_u \mu(s, k, l)_{u-1}$ .
    Save candidate list (including alignment paths) for ~

25  if user corrects result for utterance  $v$ , [ $1 \leq v \leq u$ ] then

    Retrieve utterance  $v$ 
    Retrieve candidate list and alignment paths for  $v$ 
    Identify subsets  $I_{top-choice}$  and  $C_{top-choice}$ 
30  Compute  $\Delta_v \mu(s, k, l)_u =$  for all  $s, k, l$  specified by alignment paths for all pairs in  $I_{top-choice}$  and  $C_{top-choice}$ 
    Update  $\mu(s, k, l)_u = \mu(s, k, l)_u - \omega' \Delta_v \mu(s, k, l)_u$ 

    Identify subsets  $I_{corrected}$  and  $C_{corrected}$ 
35  Compute  $\Delta_v \mu(s, k, l)_u$  for all  $s, k, l$  specified by alignment paths for all pairs in  $I_{corrected}$  and  $C_{corrected}$ 
    Update  $\mu(s, k, l)_u = \mu(s, k, l)_u - \omega' \Delta_v \mu(s, k, l)_u$ 

end

```

-21-

end

As indicated previously, adaptation is performed as recognition is performed on each utterance based on the assumption that the recognition is correct and a re-adjustment is performed only when the user makes a correction, even though that correction may be made after the user has issued several intervening utterances. This procedure is illustrated in the flowchart of Fig. 6. After a candidate list is obtained as indicated at block 151, correct (C) and incorrect (I) subsets are identified as indicated at block 153. Corrections to model parameters are computed for all pairs C and I as indicated at block 155 and the corrections are added to the then current model parameters, as indicated at block 157, using a relatively low weight. The candidate list and alignment paths are stored as indicated at block 159. If the user does not make a correction, the utterance path is incremented, as indicated at block 163, and, if there are no pending utterances, as tested at block 165, the procedure returns to the initial point to await a new utterance.

If, at the block 161, the user corrects an earlier utterance, the stored data corresponding to the item to be corrected is retrieved as indicated at block 171. Likewise, the candidate set and alignment paths for the utterance to be corrected are retrieved as indicated at block 173. The correct and incorrect subsets are identified as indicated at block 175 and the correction term is computed for all pairs in I and C as indicated at block 179. This information can either be computed at the time of correction or stored from the initial recognition. The corresponding correction factor is subtracted from the then extant model

-22-

parameters without attempting to undue all intervening corrections which may have been applied. The subset for the C (after correction) and I (after correction) are identified as indicated at block 183 and correction terms are computed for all pairs in I and C as indicated at block 183. This correction factor is then added to  
5 the model parameters as indicated at block 191. As indicated, a relatively low weighting factor is used in this adjustment since it is based on a single example rather than a batch of examples as was the case of the adjustments made during the initial or batch training illustrated in Fig. 5.

10 In view of the foregoing it may be seen that several objects of the present invention are achieved and other advantageous results have been attained.

As various changes could be made in the above constructions without departing from the scope of the invention, it should be understood that all matter contained in the above description or shown in the accompanying drawings shall  
15 be interpreted as illustrative and not in a limiting sense.

-23-

CLAIMS

What is claimed is:

1. In a speech recognition system in which text is created from a sequence of input utterances and in which each input utterance is converted to a sequence of  
5 raw vectors, a method of identifying, from a vocabulary of word models, those models which best correspond to said utterance and for improving the word models based on corrections made by the user, said method comprising:  
  
providing, for each word model state, a continuous probability distribution  
10 function (PDF) characterized by preselected parameters;  
  
using the respective continuous PDFs, calculating a ranked scoring of a group of word models; and  
  
15 inserting the word corresponding to the best scoring model into the text being created;  
  
after each utterance, adjusting said preselected parameters to increase the difference in scoring between the top and next ranking models;  
  
20 storing said raw vectors for each of a predetermined plurality of successive utterances and storing the identities of the word models in the respective selected group thereby creating a historical record of recognition events;

-24-

if a user corrects a prior recognition event in said text by selecting a different word model from the respective selected group

- 1) performing, on the current state of said parameters, an adjustment opposite to that performed with the original recognition event
- 2) performing, on the then current state of said parameters, an adjustment equivalent to that which would have been performed if said different word model had been best scoring.

2. A method as set forth in claim 1 further comprising:

identifying, for each raw vector in an input utterance, that one of a preselected plurality of quantized vectors which best matches the raw vector;

providing, for each model, a sequence of states representing the corresponding vocabulary word, said states being selected from a preselected group of states;

providing, for each word model state, both a discrete probability distribution function (PDF) and said continuous PDF;

20

providing a stored table containing distance degree-of-match metric values for each combination of quantized vector with model state as characterized by said discrete PDFs;

-25-

using the respective discrete PDFs, performing an alignment of at least selected ones of said word models with an input utterance and generating initial match scores using said stored table and said alignments;

- 5 selecting a group of well matching word models based on said initial match scores for further scoring based on said continuous PDFs.

3. A method as set forth in claim 2 wherein the scoring based on said continuous PDFs follows the alignment derived from said discrete PDFs.

10

4. In a speech recognition system in which text is created from a sequence of input utterances and in which each input utterance is converted to a sequence of raw vectors, a method of identifying, from a vocabulary of word models, those models which best correspond (match) to said utterance, said method comprising:

15

identifying, for each raw vector in an input utterance, that one of a preselected plurality of quantized vectors which best matches the raw vector;

providing, for each model, a sequence of states representing the corresponding

20 vocabulary word, said states being selected from a preselected group of states;

providing, for each word model state, both a discrete probability distribution function (PDF) and a continuous PDF characterized by Preselected parameters;

-26-

providing a stored table containing distance degree-of-match metric values for each combination of quantized vector with model state as characterized by said discrete PDFs;

using the respective discrete PDFs, performing an alignment of at least selected

5 ones of said word models with an input utterance and generating initial match scores using said stored table and said alignments;

selecting a group of well matching word models based on said initial match scores;

10

using the respective continuous PDFs and following the alignments obtained using said discrete PDFs, calculating a ranked scoring of said group of word models; and

15 inserting the word corresponding to the best scoring model into the text being created;

after each utterance, adjusting said preselected parameters to increase (by a small proportion) the difference in scoring between the top and next ranking models.

20 5. A method as set forth in claim 4 further comprising:

storing said raw vectors for each of a predetermined plurality of successive utterances and storing the identities of the word models in the respective selected

-27-

group thereby creating a historical record of recognition events;

if a user corrects a prior recognition event in said text by selecting a different word model from the respective selected group

5

1) performing, on the current state of said parameters, an adjustment opposite to that performed with the original recognition event

10

2) performing, on the then current state of said parameters, an adjustment equivalent to that which would have been performed if said different word model had been best scoring.

59385

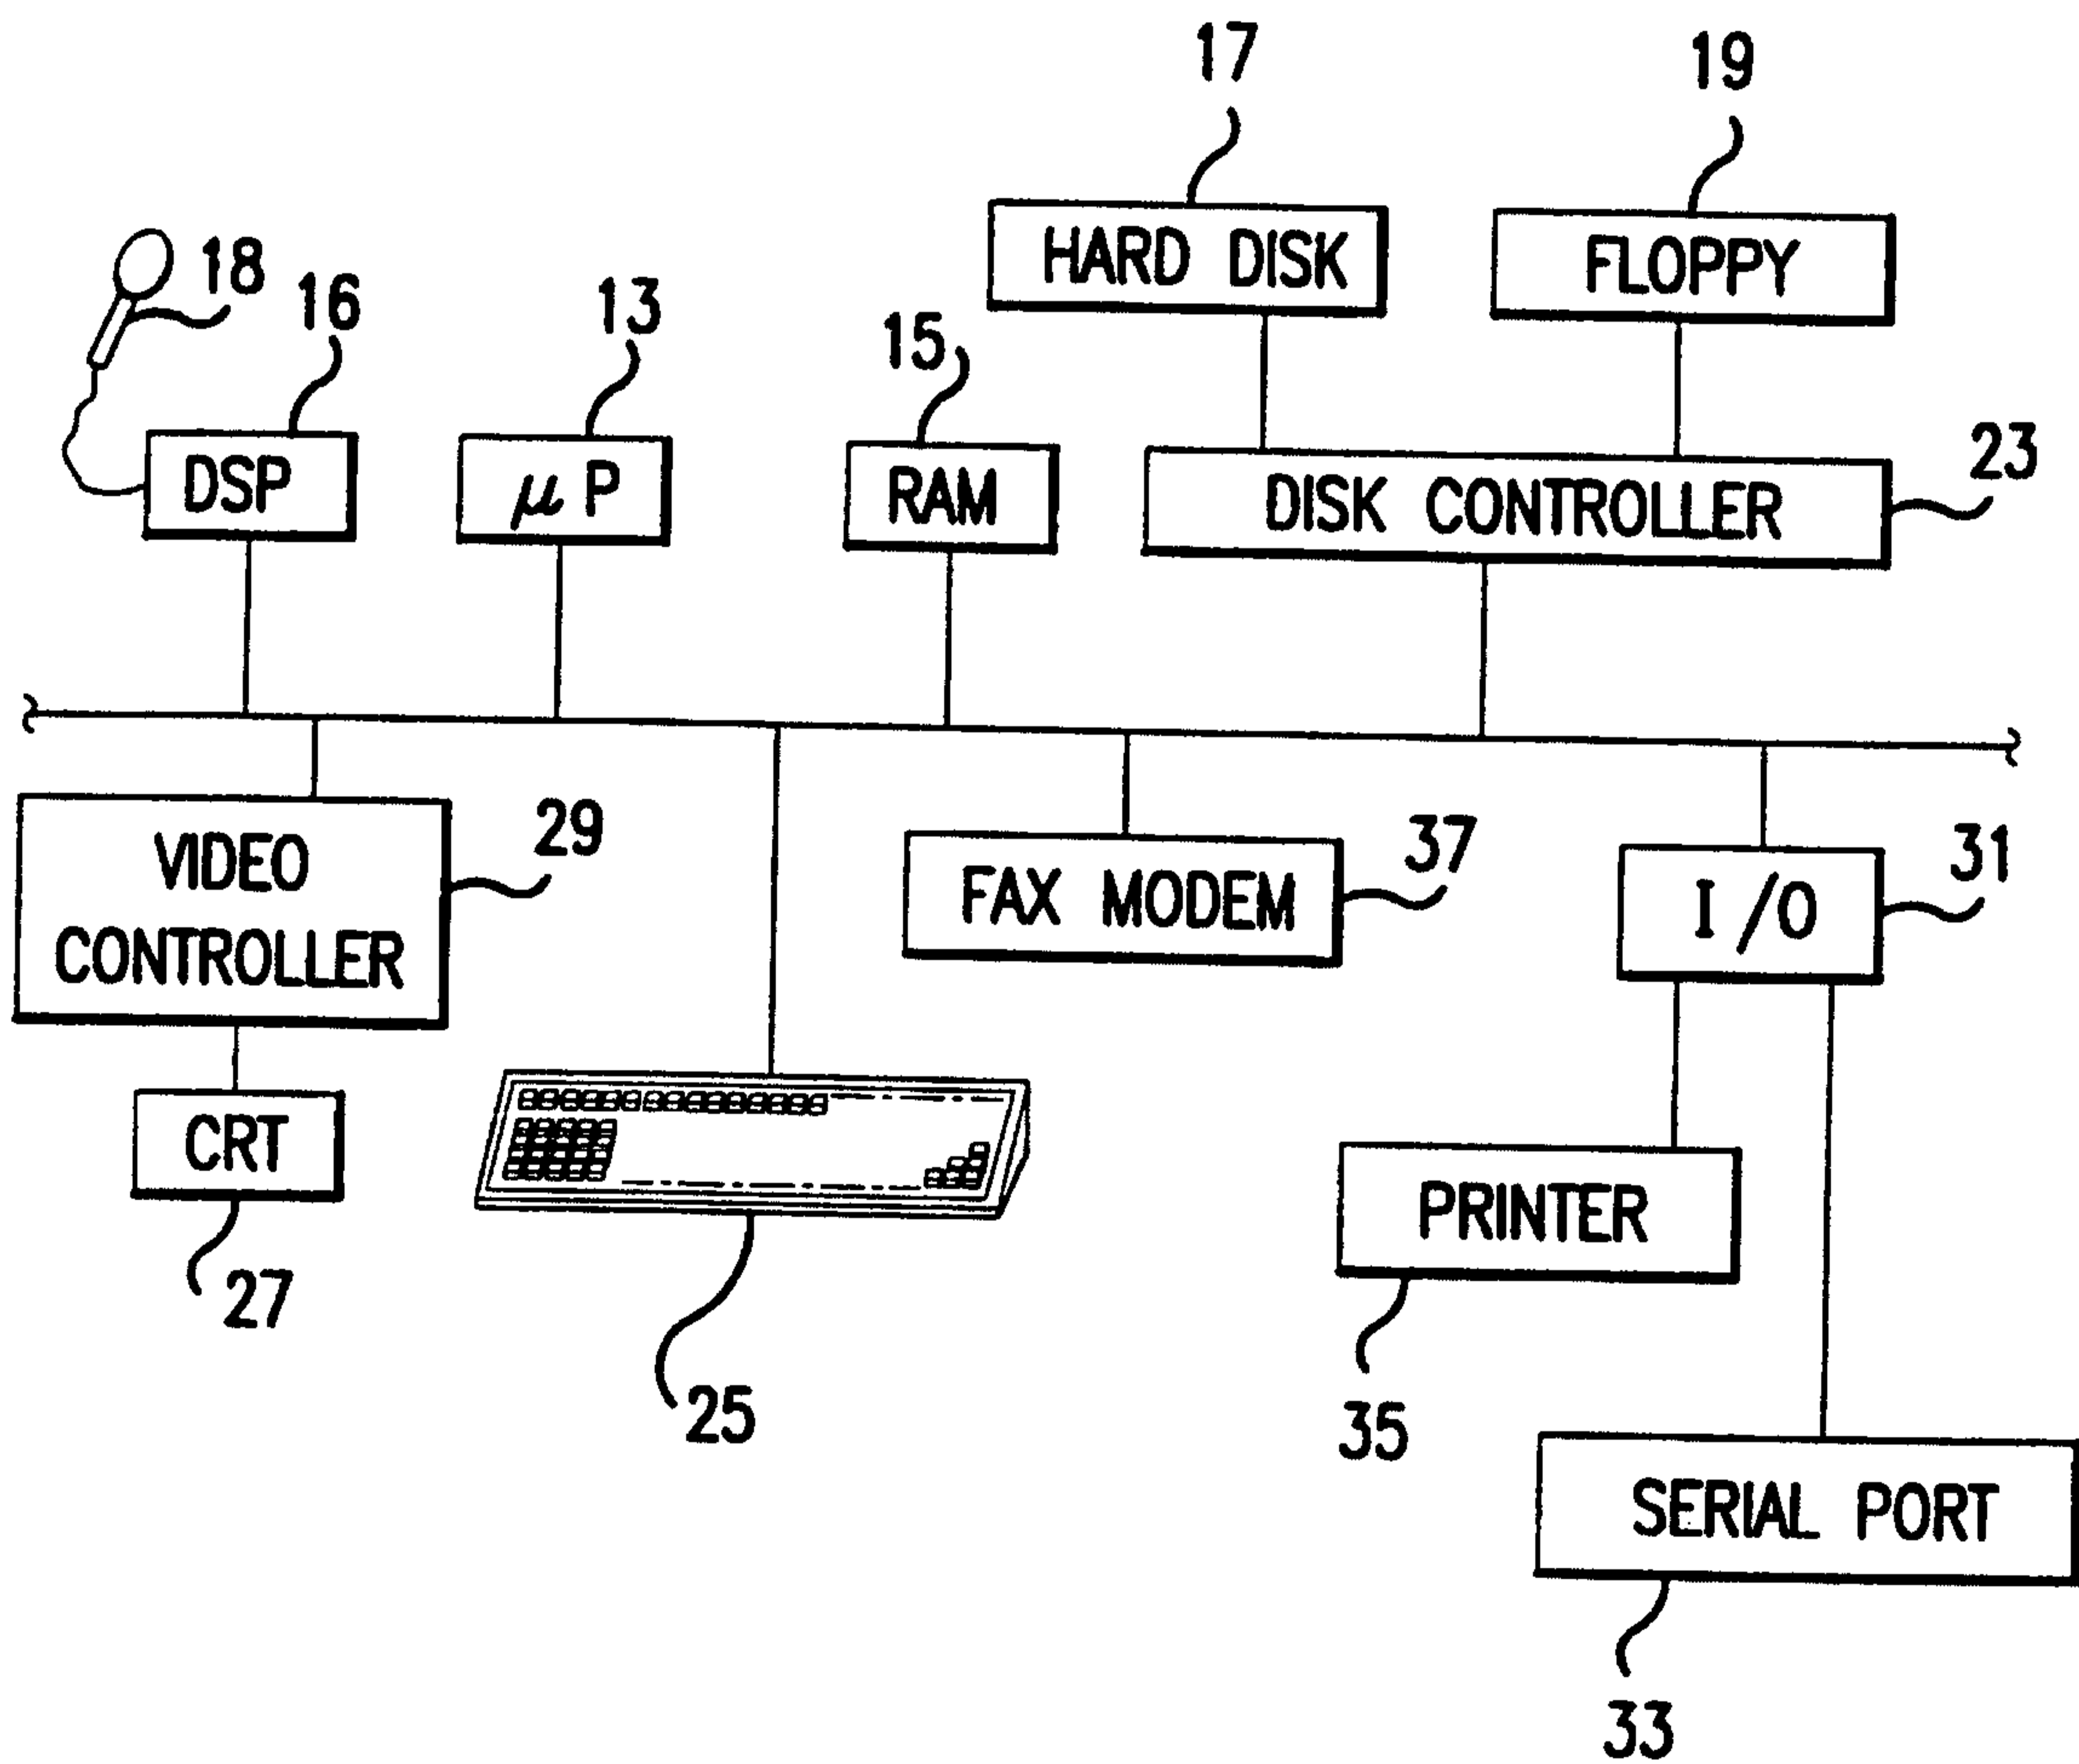


FIG. 1

2/5

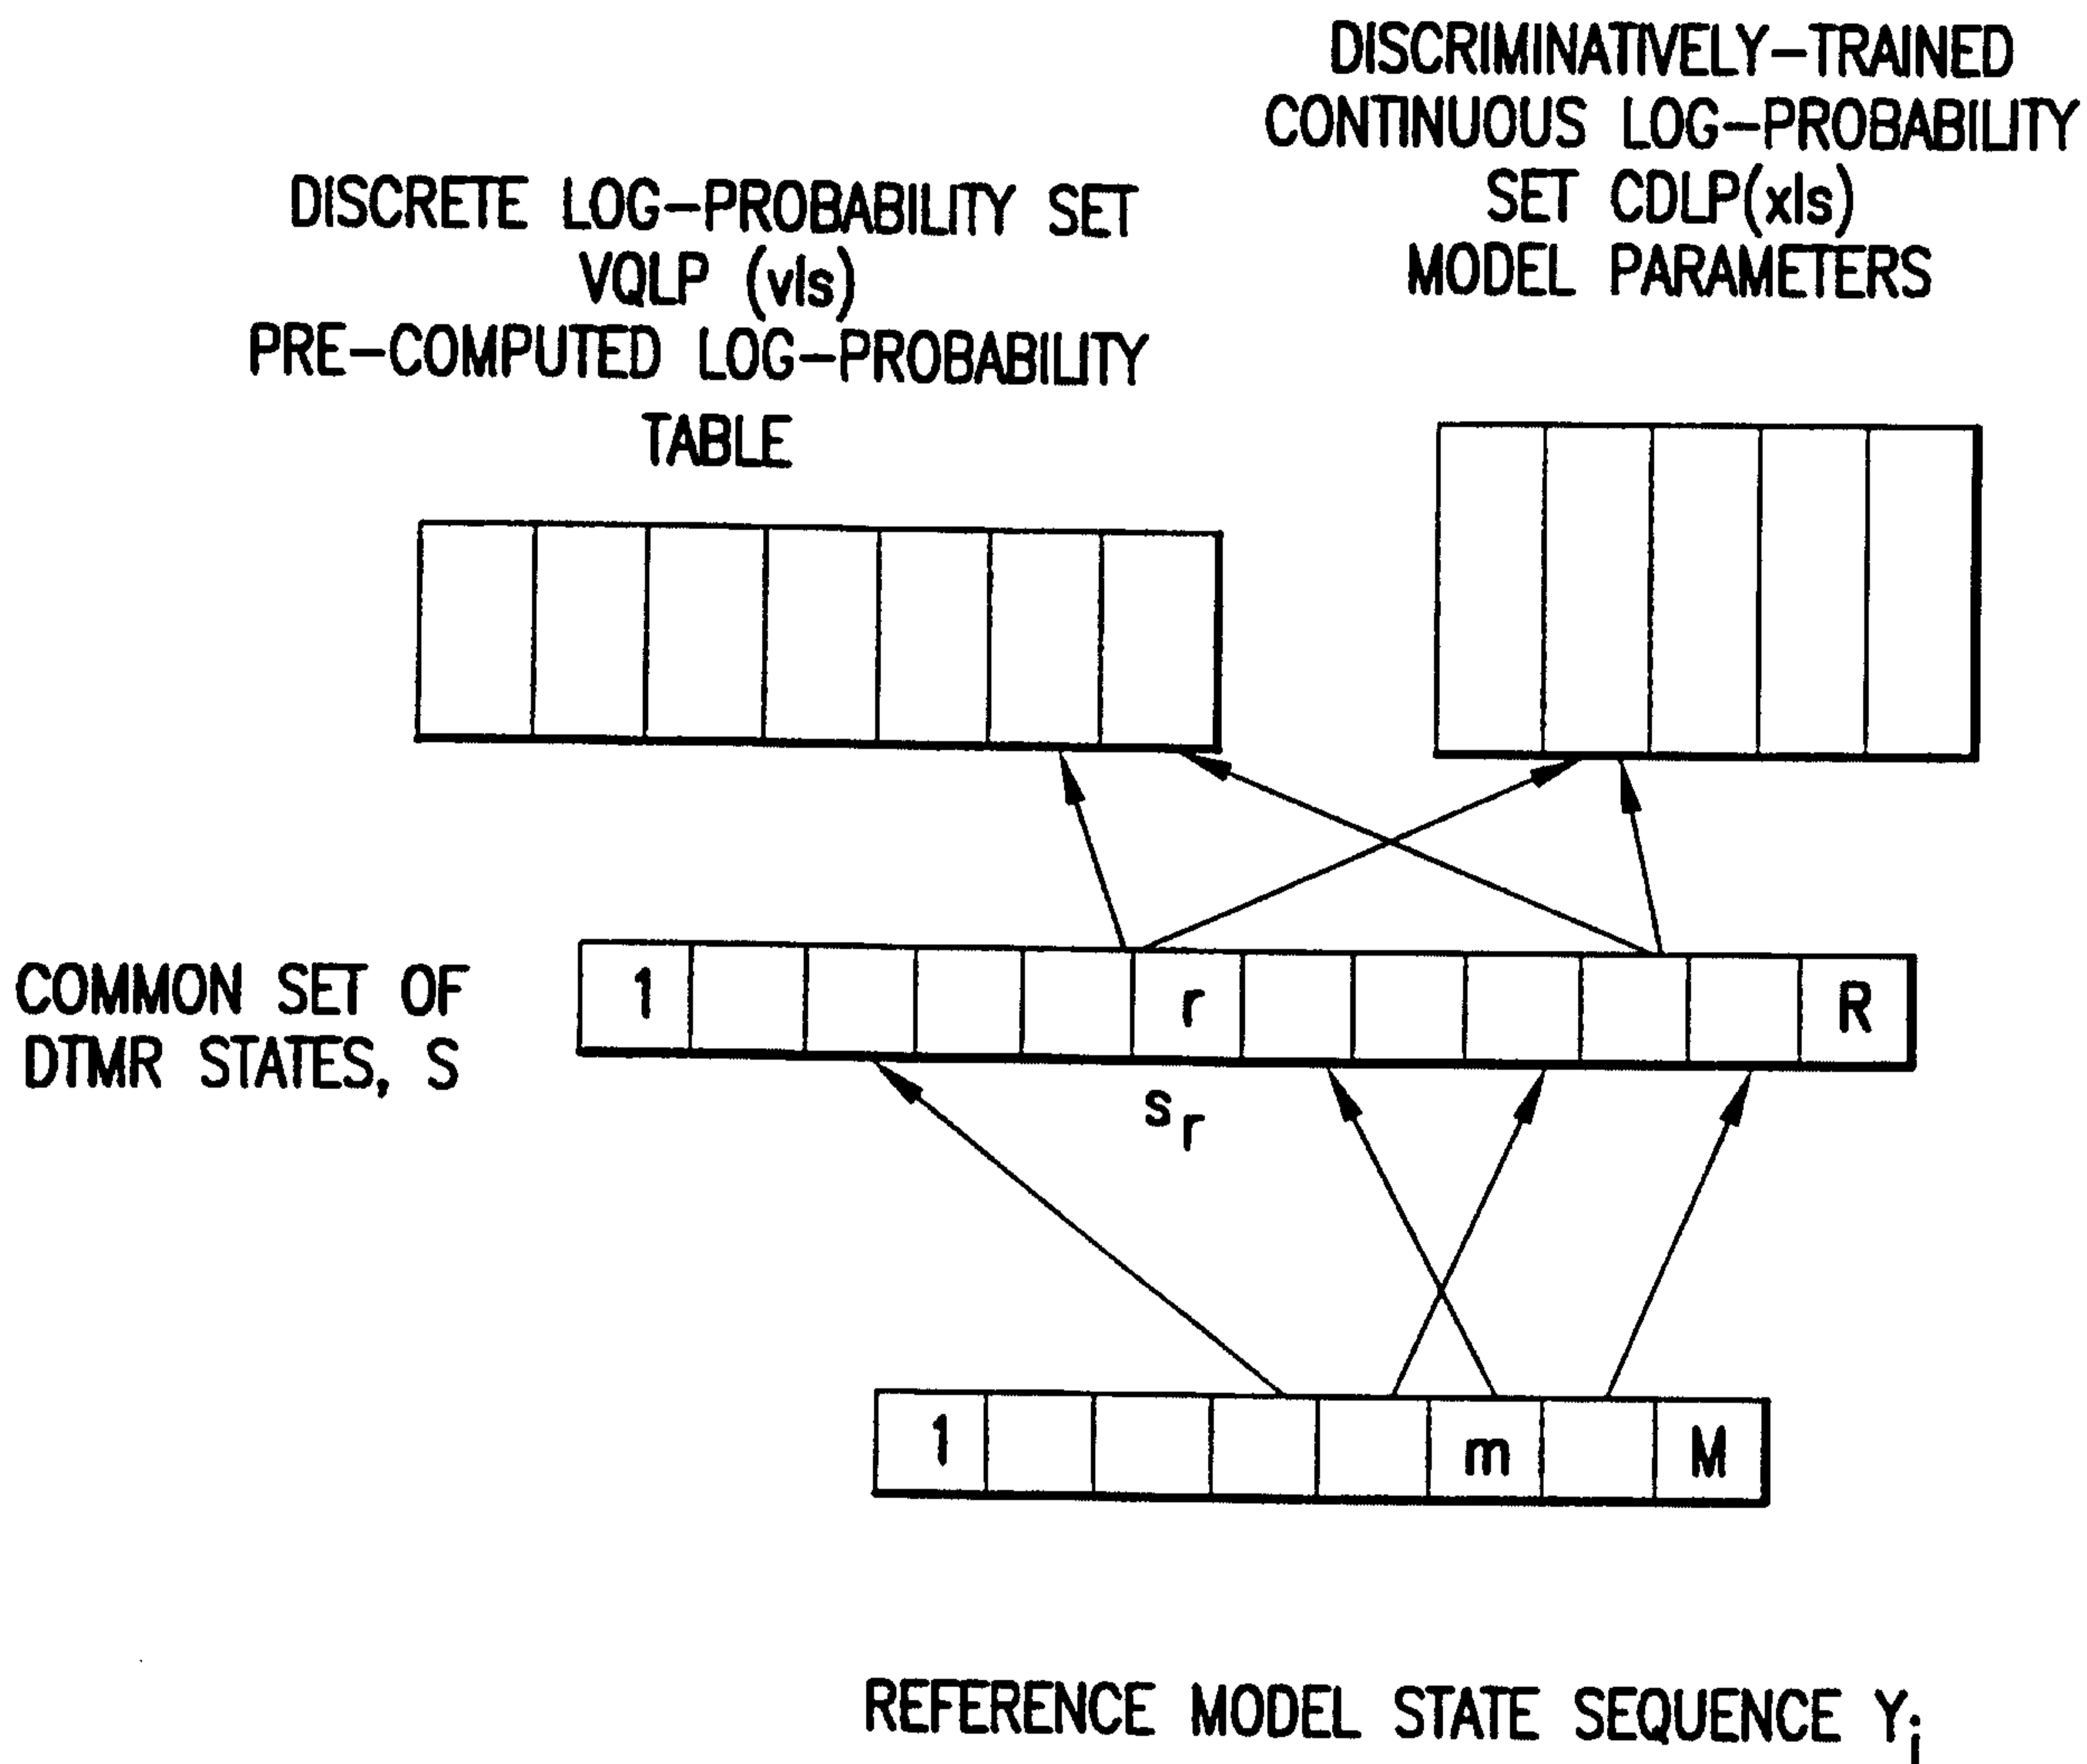


FIG.2

3/5

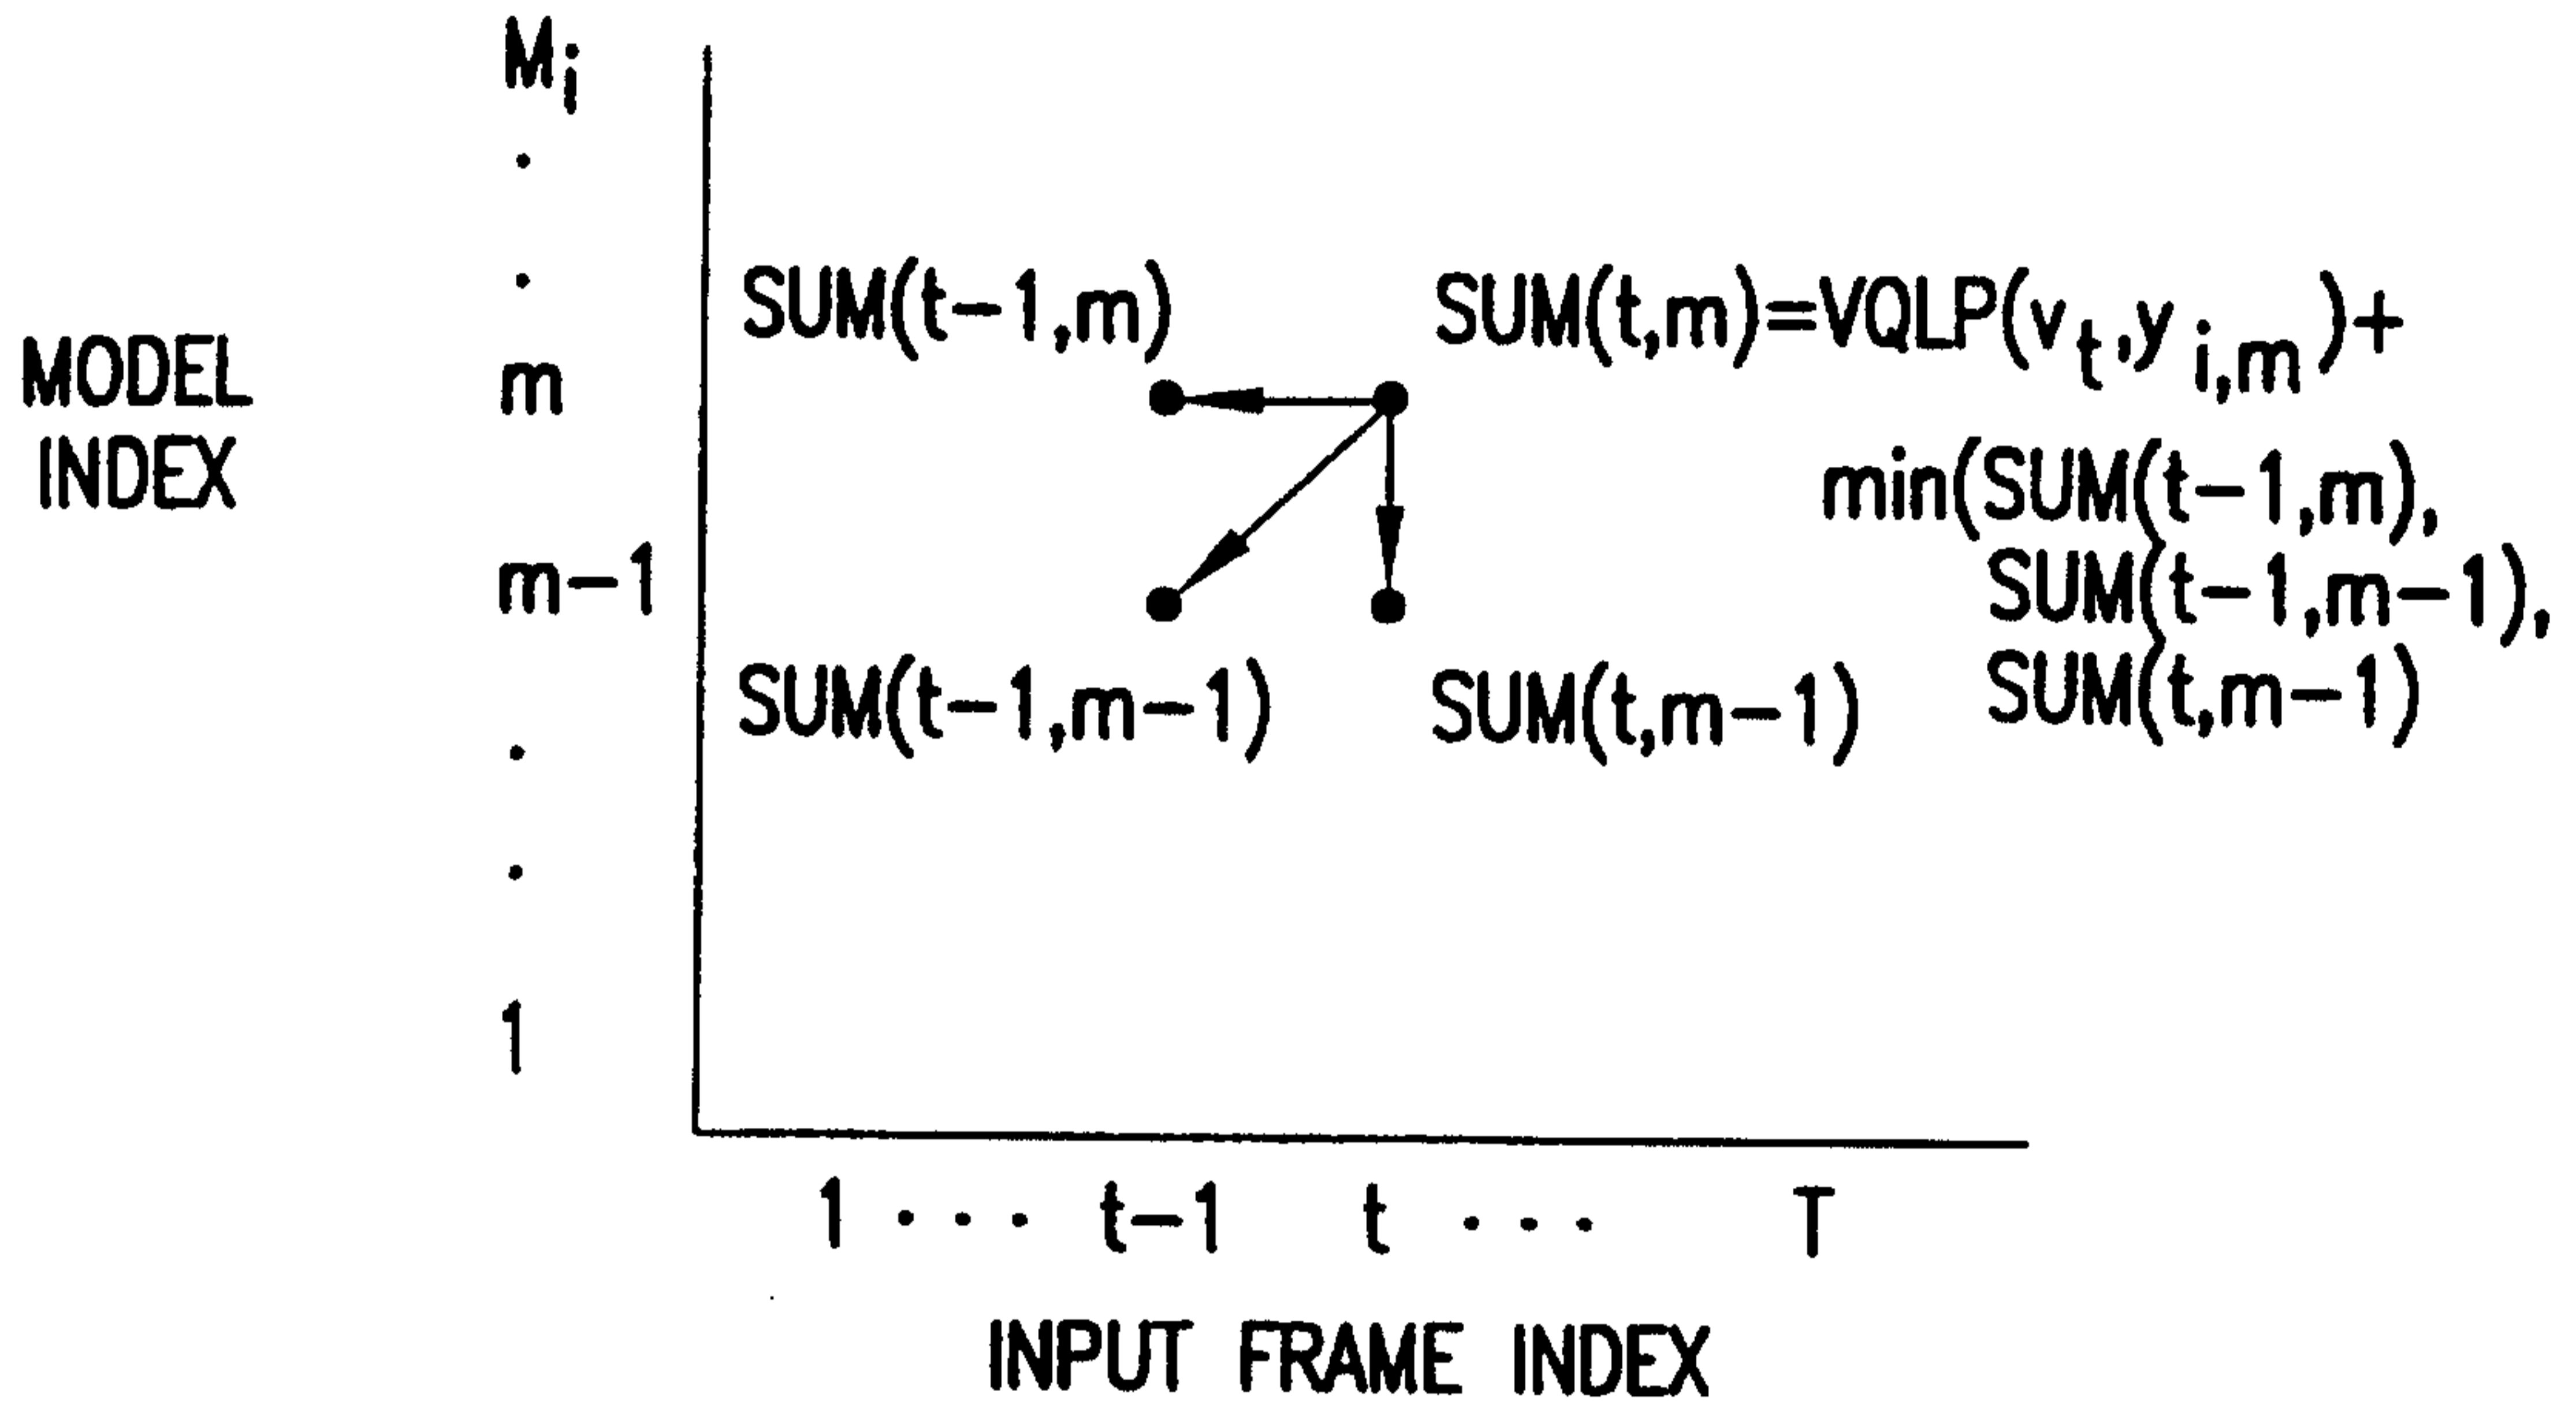


FIG.3

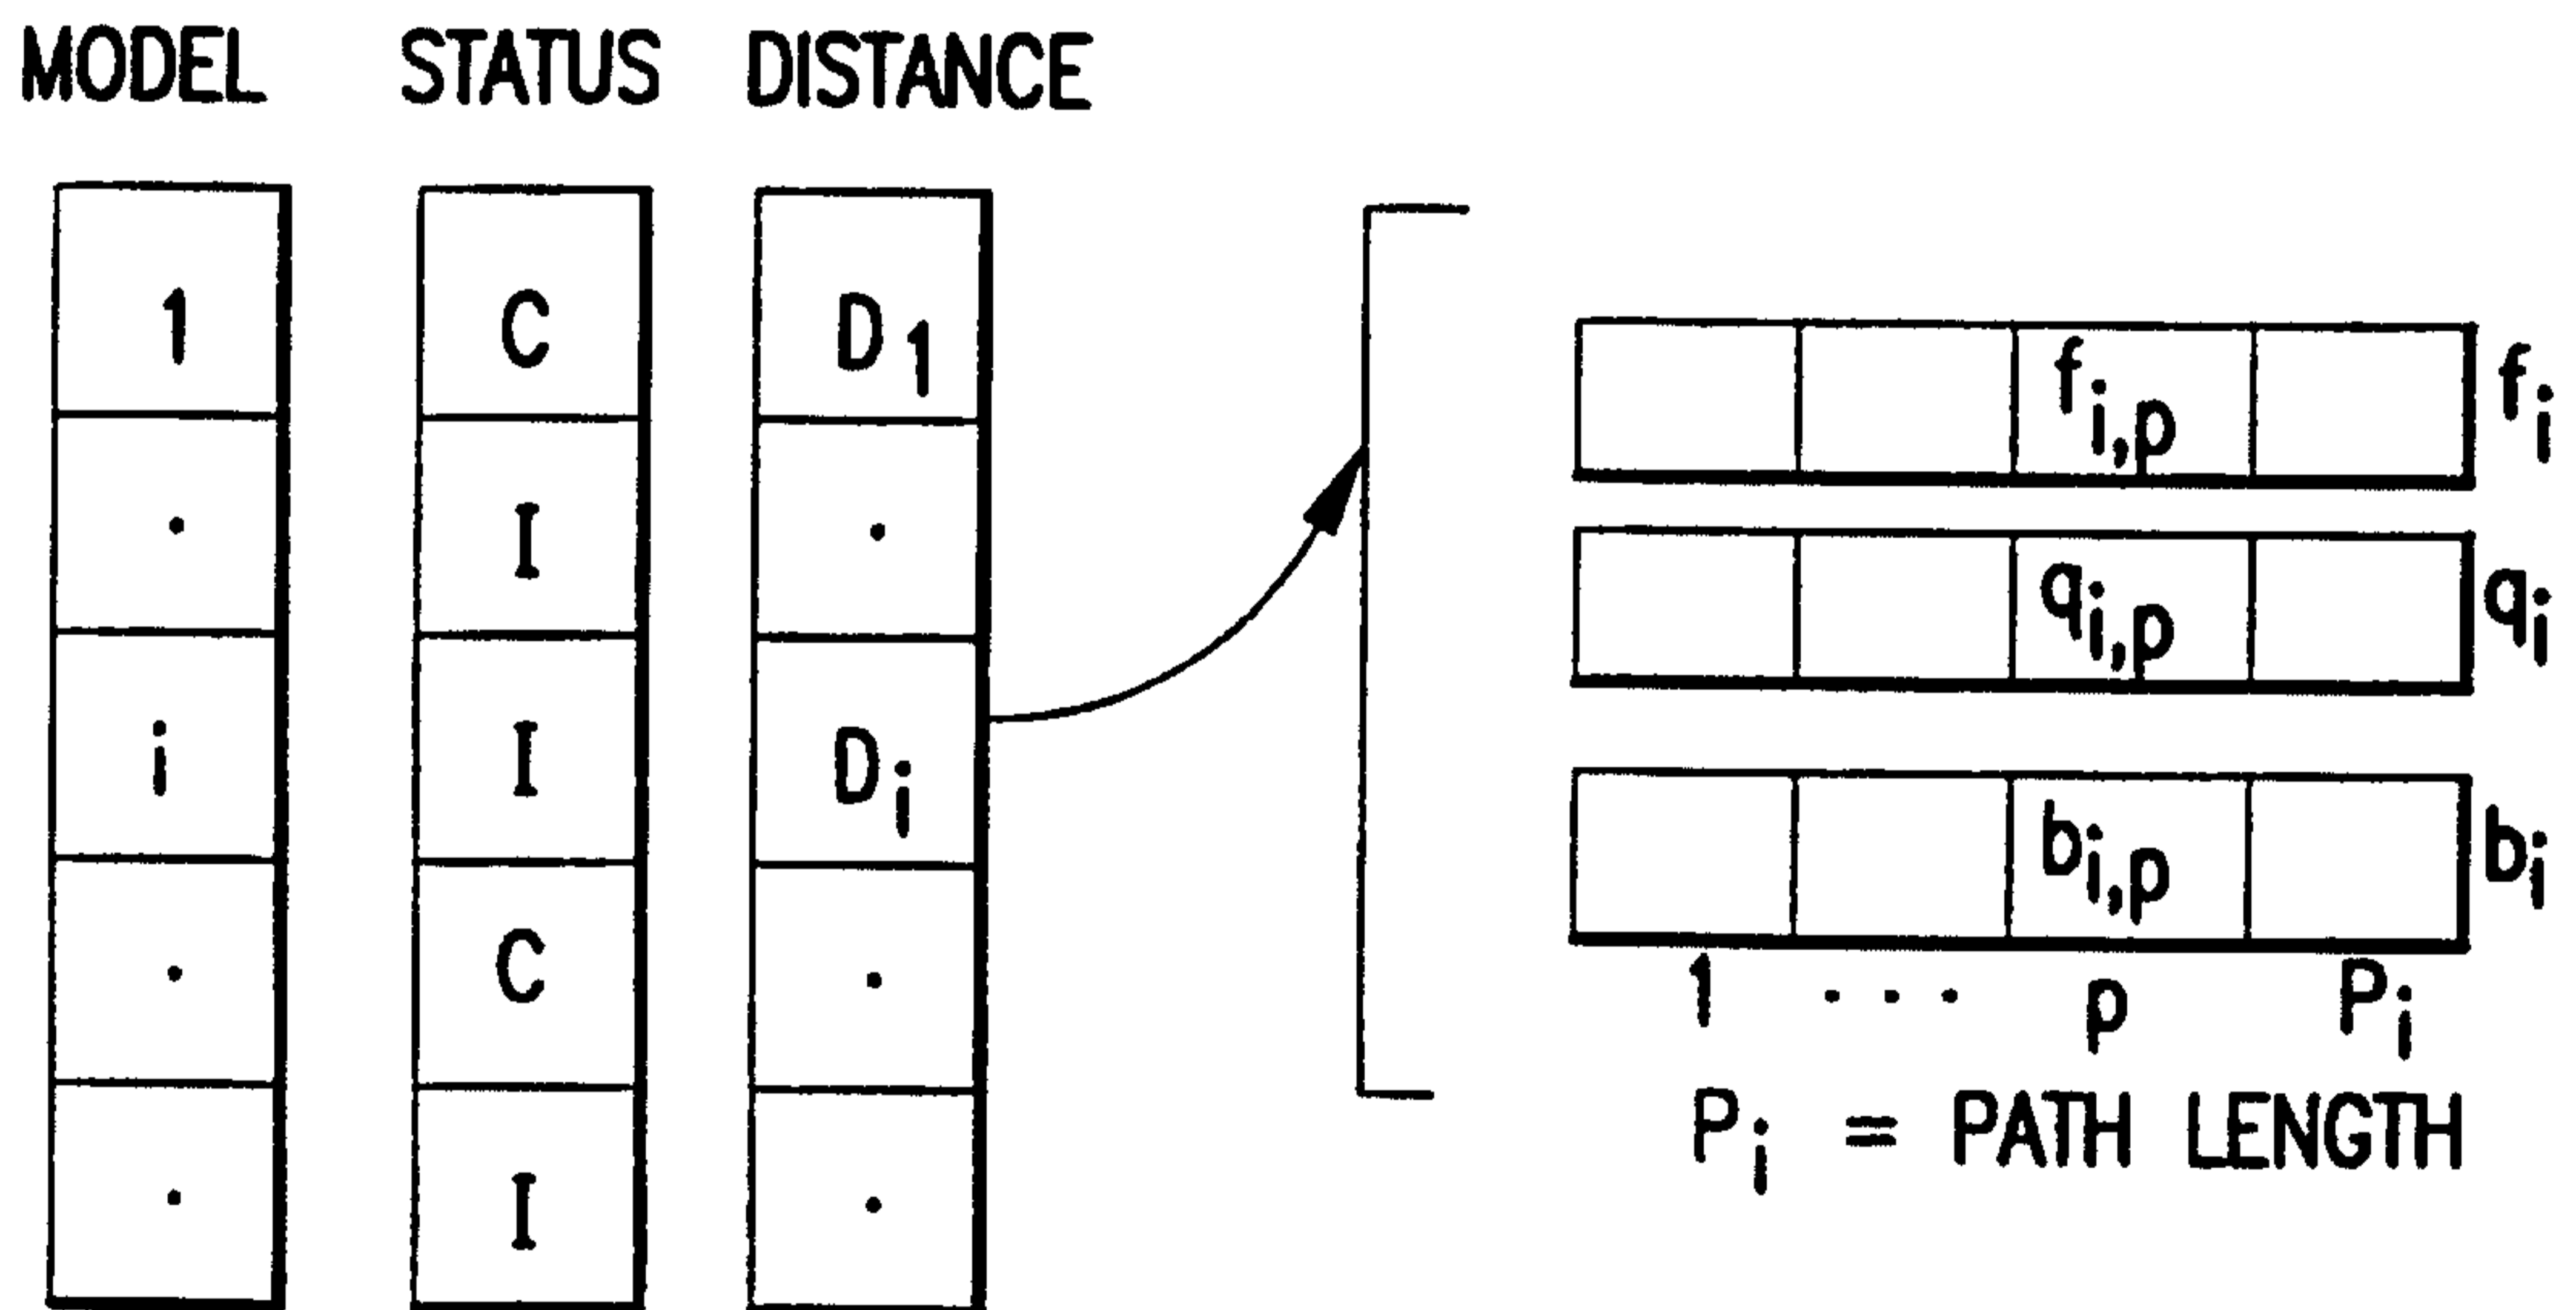


FIG.4

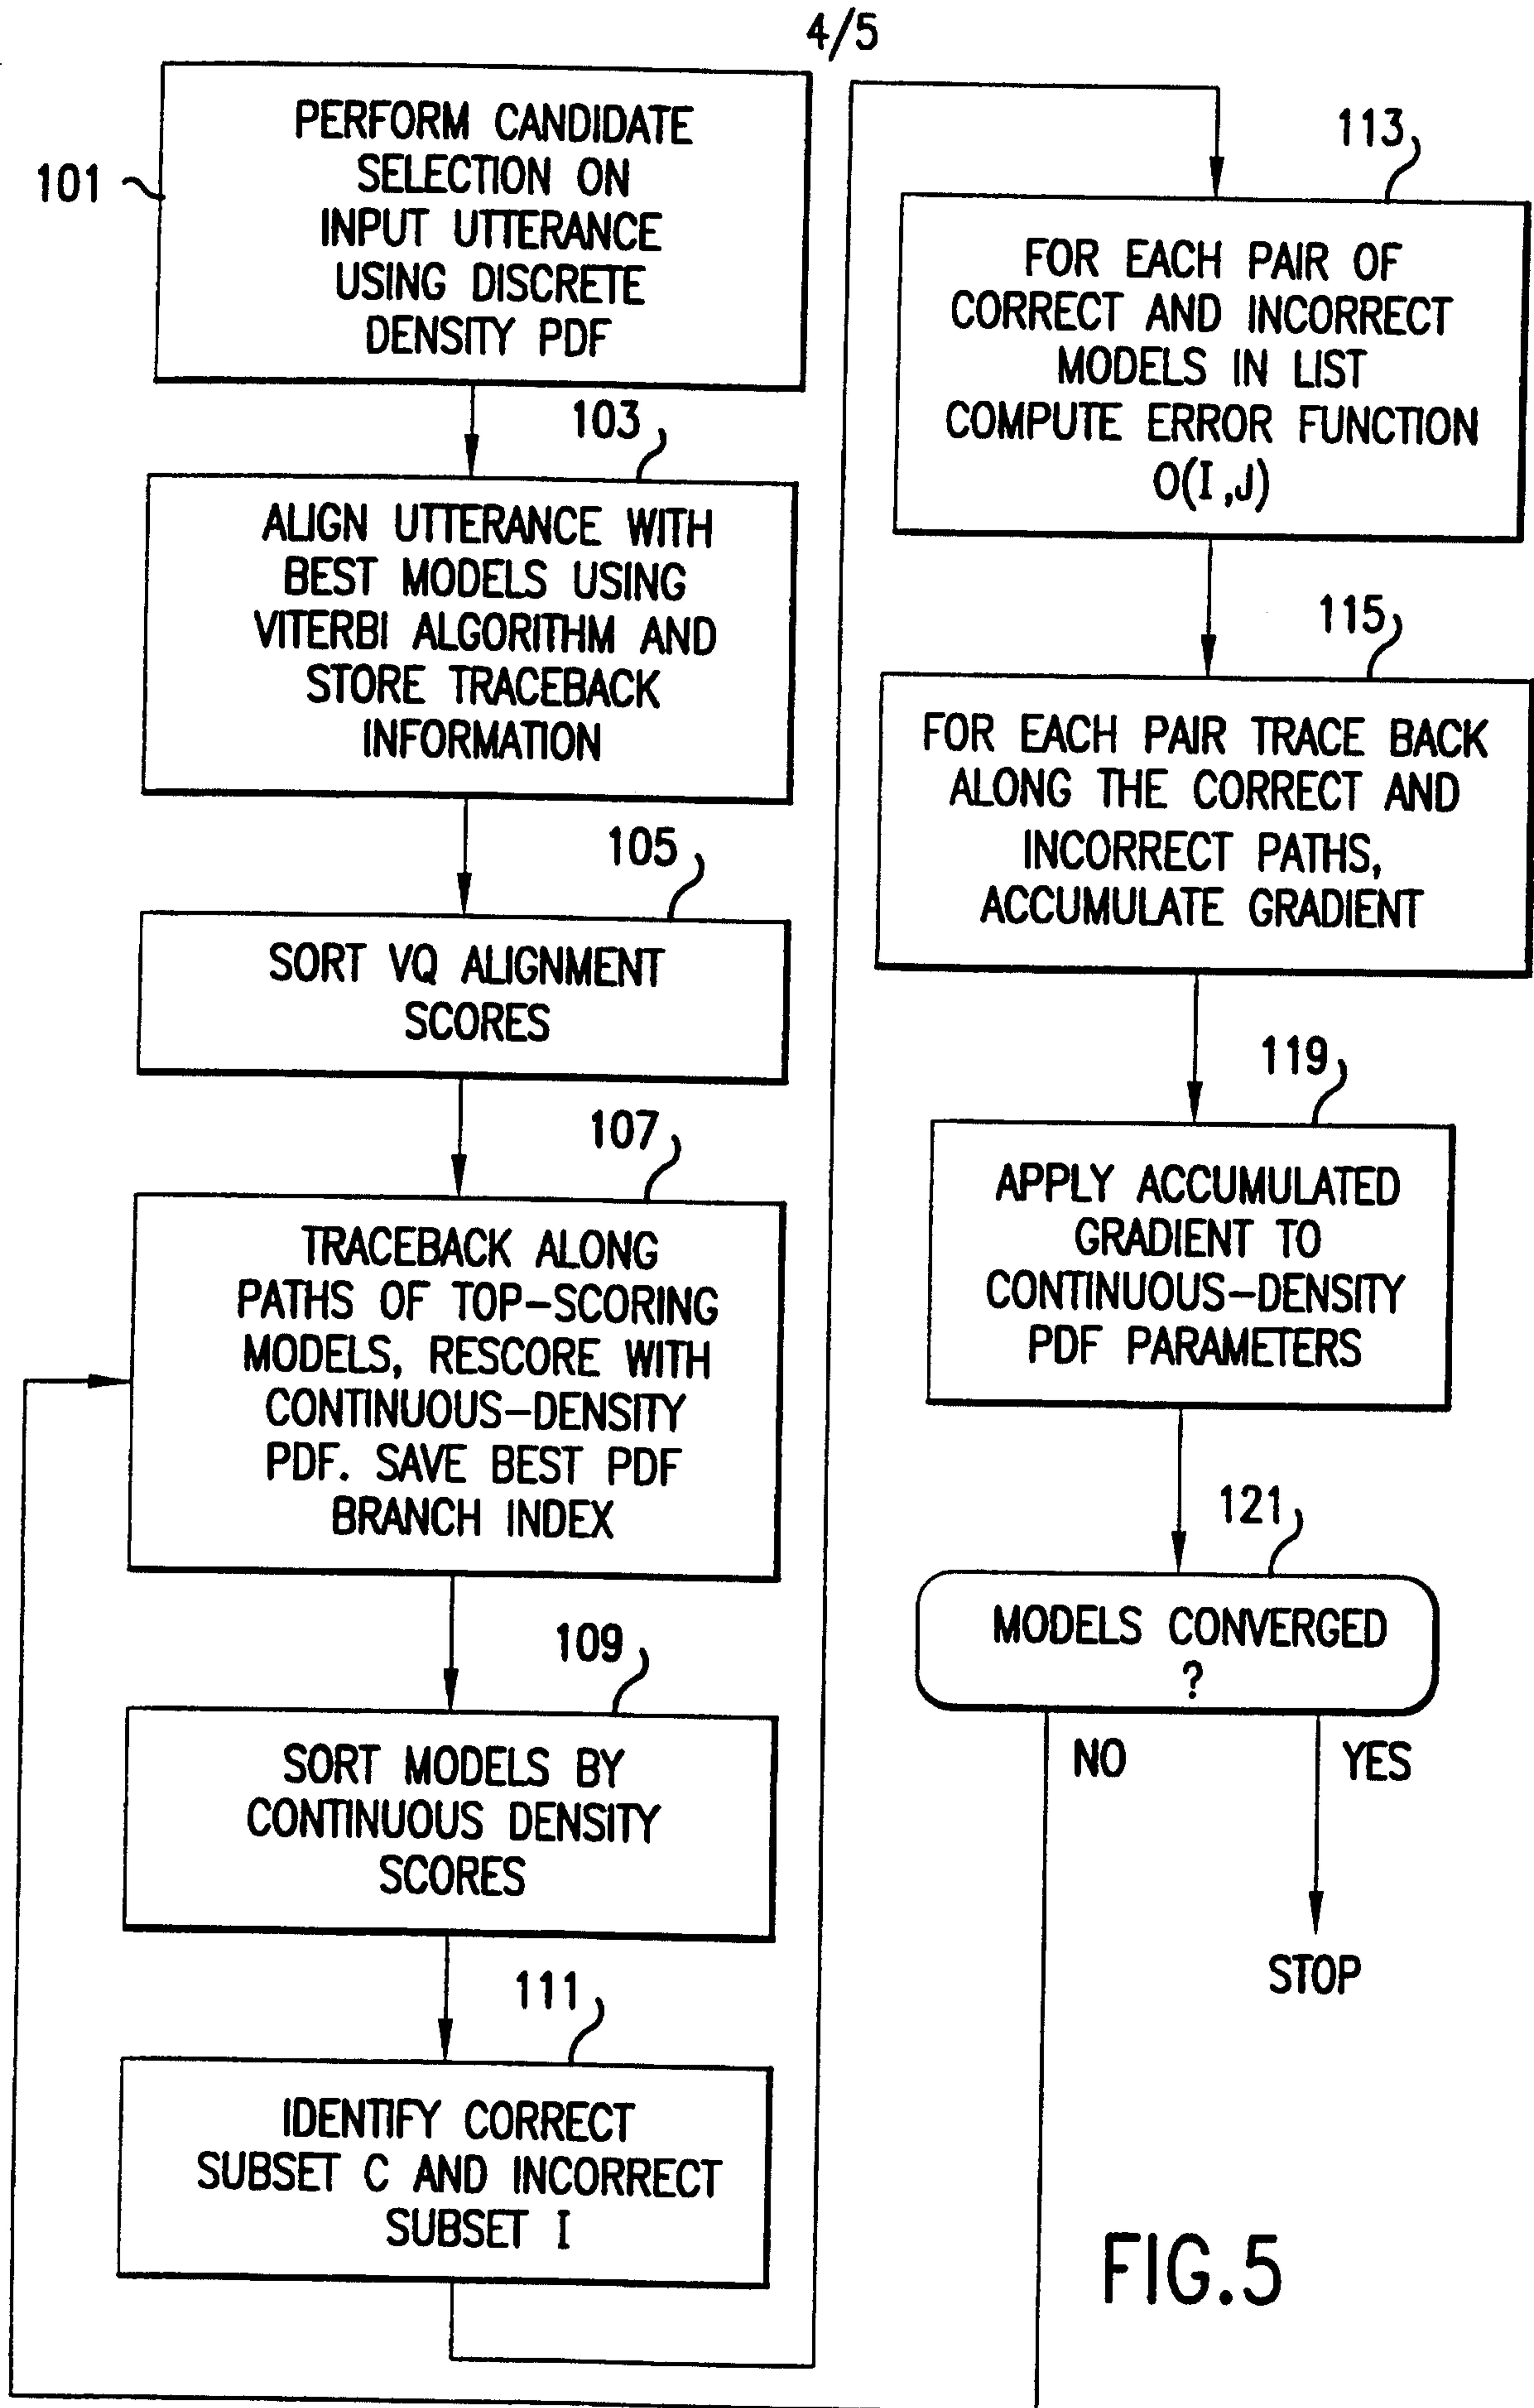


FIG.5

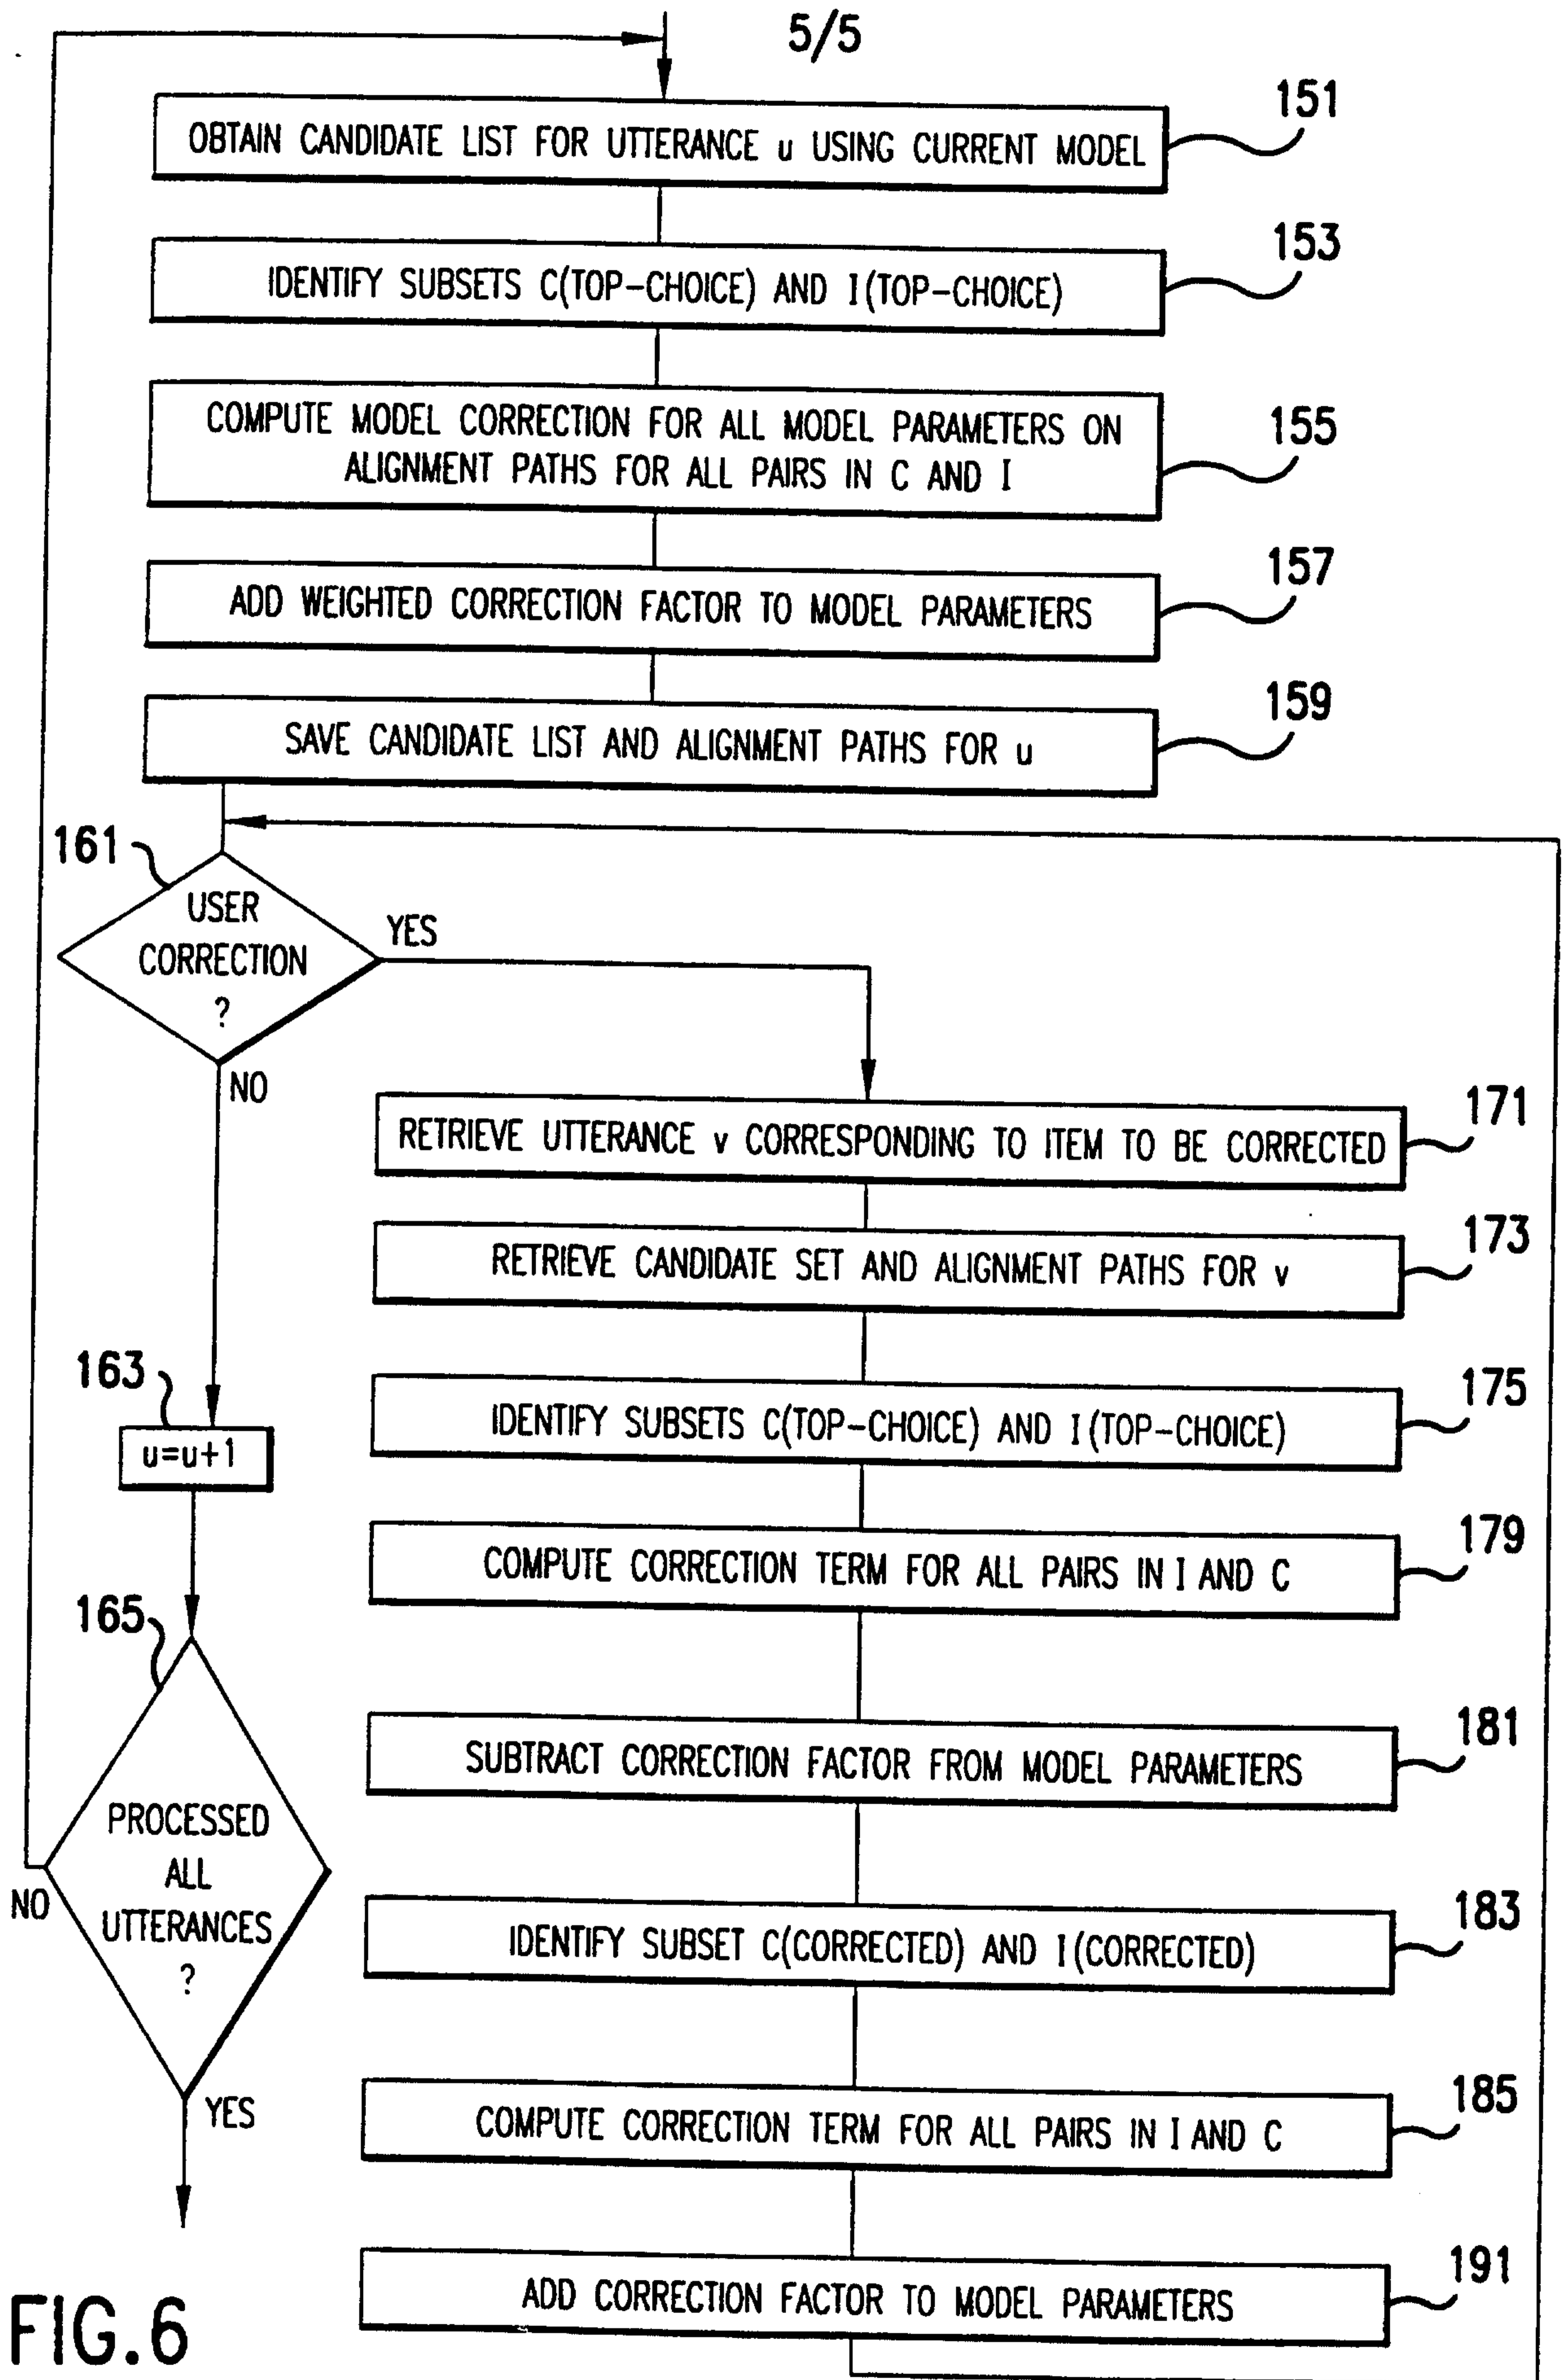


FIG. 6