

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 December 2003 (04.12.2003)


PCT

(10) International Publication Number
WO 03/100769 A1

- (51) International Patent Classification⁷: **G10L 15/20**, 15/02, 21/02
- (21) International Application Number: PCT/US03/16032
- (22) International Filing Date: 20 May 2003 (20.05.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/152,127 20 May 2002 (20.05.2002) US
10/152,143 20 May 2002 (20.05.2002) US
10/236,042 5 September 2002 (05.09.2002) US
- (71) Applicant: **MICROSOFT CORPORATION** [US/US];
One Microsoft Way, Redmond, WA 98052 (US).
- (72) Inventors: **DENG, Li**; 22310 N.E. 30th Street, Sammamish, WA 98074 (US). **ACERO, Alejandro**; 6525 163rd Place S.E., Bellevue, WA 98006 (US). **DROPPA, James, G.**; 13820 282nd Court, N.E., Duvall, WA 98019 (US).
- (74) Agents: **MAGEE, Theodore, M.** et al.; Westman, Champlin & Kelly, P.A., Suite 1600, International Centre, 900 Second Avenue South, Minneapolis, MN 55402-3319 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT (utility model), AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK (utility model), SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report

[Continued on next page]

(54) Title: METHOD OF DETERMINING UNCERTAINTY ASSOCIATED WITH NOISE REDUCTION



(57) Abstract: A method and apparatus are provided for determining uncertainty in noise reduction based on a parametric model of speech distortion. The method is first used to reduce noise in a noisy signal. In particular, noise is reduced (304) from a representation of a portion of a noisy signal to produce a representation of a cleaned signal by utilizing an acoustic environment model (413). The uncertainty associated with the noise reduction process is then computed. In one embodiment, the uncertainty of the noise reduction process is used, in conjunction with the noise-reduced signal, to decode (306) a pattern state.



WO 03/100769 A1



-
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

METHOD OF DETERMINING UNCERTAINTY ASSOCIATED WITH NOISE REDUCTION

BACKGROUND OF THE INVENTION

5 The present invention relates to pattern recognition. In particular, the present invention relates to performing pattern recognition after noise reduction.

 A pattern recognition system, such as a
10 speech recognition system, takes an input signal and attempts to decode the signal to find a pattern represented by the signal. For example, in a speech recognition system, a speech signal (often referred to as a test signal) is received by the recognition
15 system and is decoded to identify a string of words represented by the speech signal.

 To decode the incoming test signal, most recognition systems utilize one or more models that describe the likelihood that a portion of the test
20 signal represents a particular pattern. Examples of such models include Neural Nets, Dynamic Time Warping, segment models, and Hidden Markov Models.

 Before a model can be used to decode an incoming signal, it must be trained. This is
25 typically done by measuring input training signals generated from a known training pattern. For example, in speech recognition, a collection of speech signals is generated by speakers reading from a known text. These speech signals are then used to
30 train the models.

-2-

In order for a model to work optimally, the signals used to train the model should be similar to the eventual test signals that are decoded. In particular, it is desirable that the training signals
5 contain the same amount and type of noise as the test signals that are decoded.

Typically, the training signal is collected under "clean" conditions and is considered to be relatively noise free. To achieve this same low
10 level of noise in the test signal, many prior art systems apply noise reduction techniques to the testing data. These noise reduction techniques result in a cleaned test signal that is then used during pattern recognition. In most systems, the noise
15 reduction technique produces a sequence of multi-dimensional feature vectors, with each feature vector representing a frame of a noise-reduced signal.

Unfortunately, noise reduction techniques do not work perfectly and as a result, there is some
20 inherent uncertainty in the cleaned signal. In the past, there have been two general techniques for dealing with such uncertainty. The first has been to ignore the uncertainty and treat the noise reduction process as being perfect. Since this ignores the
25 true state of the recognition system, it results in recognition errors that could be avoided.

The other prior art technique for dealing with uncertainty in noise reduction is to identify frames of the input signal where the noise reduction
30 technique is likely to have performed poorly. In

-3-

these frames, dimensions of the feature vectors that are likely in error are marked by the noise reduction system so that they are not used during recognition. Thus, the feature vector components that have more
5 than a predetermined amount of uncertainty are completely ignored during decoding. Although such systems acknowledge uncertainty in noise reduction, the technique of completely ignoring a component treats the component as providing no information that
10 would be helpful during recognition. This is highly unlikely because even with a significant amount of uncertainty, the noise-reduced component still provides some information that would be helpful during recognition.

15 In addition, the prior art has not provided a means for determining the uncertainty of some noise-removal processes. As a result, it has not been possible to determine the uncertainty associated with those processes.

20 In light of this, techniques are needed to identify the uncertainty in noise reduction and use that uncertainty during pattern recognition.

SUMMARY OF THE INVENTION

A method and apparatus are provided for
25 reducing noise in a noisy signal. In particular, noise is reduced from a representation of a portion of a noisy signal to produce a representation of a cleaned signal by utilizing an acoustic environment model. The uncertainty associated with the noise
30 reduction process is then computed. In one

-4-

embodiment, the uncertainty of the noise reduction process is used to decode a pattern state from the cleaned signal.

BRIEF DESCRIPTION OF THE DRAWINGS

5 FIG. 1 is a block diagram of one computing environment in which the present invention may be practiced.

 FIG. 2 is a block diagram of an alternative computing environment in which the present invention
10 may be practiced.

 FIG. 3 is a flow diagram of a method of using a noise reduction system of one embodiment of the present invention.

 FIG. 4 is a block diagram of a pattern
15 recognition system in which embodiments of the present invention may be used.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

 FIG. 1 illustrates an example of a suitable computing system environment 100 on which the
20 invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the
25 computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

 The invention is operational with numerous
30 other general purpose or special purpose computing

-5-

system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal
5 computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, telephony systems, distributed computing
10 environments that include any of the above systems or devices, and the like.

The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a
15 computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed
20 computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including
25 memory storage devices.

With reference to FIG. 1, an exemplary system for implementing the invention includes a general-purpose computing device in the form of a computer 110. Components of computer 110 may
30 include, but are not limited to, a processing unit

-6-

120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

Computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 110 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape,

-7-

magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 110. Communication media
5 typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal"
10 means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired
15 connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

The system memory 130 includes computer
20 storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between
25 elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 120. By way of
30 example, and not limitation, FIG. 1 illustrates

-8-

operating system 134, application programs 135, other program modules 136, and program data 137.

The computer 110 may also include other removable/non-removable volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 141 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

The drives and their associated computer storage media discussed above and illustrated in FIG. 1, provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In FIG. 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules

-9-

146, and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137.
5 Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers here to illustrate that, at a minimum, they are different copies.

A user may enter commands and information
10 into the computer 110 through input devices such as a keyboard 162, a microphone 163, and a pointing device 161, such as a mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish, scanner, or the like.
15 These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a
20 universal serial bus (USB). A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to the monitor, computers may also include other peripheral output devices such
25 as speakers 197 and printer 196, which may be connected through an output peripheral interface 190.

The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The
30 remote computer 180 may be a personal computer, a

-10-

hand-held device, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 110. The
5 logical connections depicted in FIG. 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the
10 Internet.

When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110
15 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160, or other
20 appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 1 illustrates remote application
25 programs 185 as residing on remote computer 180. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

-11-

FIG. 2 is a block diagram of a mobile device 200, which is an exemplary computing environment. Mobile device 200 includes a microprocessor 202, memory 204, input/output (I/O) components 206, and a communication interface 208 for communicating with remote computers or other mobile devices. In one embodiment, the afore-mentioned components are coupled for communication with one another over a suitable bus 210.

Memory 204 is implemented as non-volatile electronic memory such as random access memory (RAM) with a battery back-up module (not shown) such that information stored in memory 204 is not lost when the general power to mobile device 200 is shut down. A portion of memory 204 is preferably allocated as addressable memory for program execution, while another portion of memory 204 is preferably used for storage, such as to simulate storage on a disk drive.

Memory 204 includes an operating system 212, application programs 214 as well as an object store 216. During operation, operating system 212 is preferably executed by processor 202 from memory 204. Operating system 212, in one preferred embodiment, is a WINDOWS® CE brand operating system commercially available from Microsoft Corporation. Operating system 212 is preferably designed for mobile devices, and implements database features that can be utilized by applications 214 through a set of exposed application programming interfaces and methods. The objects in object store 216 are maintained by

-12-

applications 214 and operating system 212, at least partially in response to calls to the exposed application programming interfaces and methods.

Communication interface 208 represents
5 numerous devices and technologies that allow mobile device 200 to send and receive information. The devices include wired and wireless modems, satellite receivers and broadcast tuners to name a few. Mobile device 200 can also be directly connected to a
10 computer to exchange data therewith. In such cases, communication interface 208 can be an infrared transceiver or a serial or parallel communication connection, all of which are capable of transmitting streaming information.

15 Input/output components 206 include a variety of input devices such as a touch-sensitive screen, buttons, rollers, and a microphone as well as a variety of output devices including an audio generator, a vibrating device, and a display. The
20 devices listed above are by way of example and need not all be present on mobile device 200. In addition, other input/output devices may be attached to or found with mobile device 200 within the scope of the present invention.

25 In the discussion below, the present invention is described with reference to speech recognition to facilitate understanding of the invention. However, those skilled in the art will recognize that the present invention is not limited

-13-

to speech recognition and may be used in any pattern recognition system.

The goal of speech recognition is to convert a speech signal into a sequence of words. To
5 do this, the speech signal is typically divided into overlapping frames of speech and each frame is converted into a feature vector, $\mathbf{y}$, which describes the content of the frame. During recognition, also referred to as decoding, an acoustic model is used to
10 identify a most likely sequence of phonetic states that could be represented by the feature vectors. The resulting sequence of phonetic states represents the decoded sequence of words.

Many acoustic models evaluate a mixture of
15 Gaussians to determine the probability of any one phonetic state. Thus, at each frame, a mixture of Gaussians is evaluated for each phonetic state to identify the most likely phonetic state for the frame. The evaluation of each Gaussian can be viewed
20 as determining $p(\mathbf{y}|m)$, which is the probability of the feature vector, $\mathbf{y}$, given the mixture component m for the phonetic state.

When using noise-reduction in combination with speech recognition, a noisy input feature
25 vector, $\mathbf{y}$, is converted into a cleaned feature vector, $\hat{\mathbf{x}}$, by the noise-reduction system. The cleaned feature vector represents an estimate of an actual clean feature vector, $\mathbf{x}$. The cleaned feature

-14-

vector, $\hat{\mathbf{x}}$, is then applied as an input signal to the speech recognition system, which attempts to identify a most-likely phonetic state given the cleaned feature vector.

5 Without noise corruption, the observation vector, $\mathbf{y}$, is the same as the clean vector, $\mathbf{x}$. As such, evaluating a Gaussian for a mixture component for a clean observation vector $p(\mathbf{y}|m)$ is the same as evaluating the probability of clean vector $\mathbf{x}$ given
10 mixture component m , $p(\mathbf{x}|m)$.

 However, when the observation vector is corrupted by noise and a noise-reduction system is used, the input to the recognizer is no longer the ideal clean vector, $\mathbf{x}$, but instead is the cleaned
15 feature vector, $\hat{\mathbf{x}}$, which is only an estimate of $\mathbf{x}$. In the past, it has been common to simply treat the cleaned feature vector, $\hat{\mathbf{x}}$, as if it were equal to the ideal clean feature vector, $\mathbf{x}$. In such cases, evaluating a Gaussian for a mixture component $p(\mathbf{y}|m)$
20 has been approximated by determining the probability $p(\hat{\mathbf{x}}|m)$.

 Under the present invention, a more rigorous approach is used in which the probability $p(\mathbf{y}|m)$ is determined by marginalizing a joint
25 conditional probability density function $p(\mathbf{y}, \mathbf{x}|m)$ over

-15-

all possible unseen clean speech vectors. In terms of an equation:

$$p(\mathbf{y}|m) = \int_{-\infty}^{\infty} p(\mathbf{y}, \mathbf{x}|m) d\mathbf{x} \quad \text{EQ. 1}$$

The right-hand side of EQ. 1 may be expanded so that EQ. 1 becomes:

$$p(\mathbf{y}|m) = \int_{-\infty}^{\infty} p(\mathbf{y}|\mathbf{x}, m) p(\mathbf{x}|m) d\mathbf{x} \quad \text{EQ. 2}$$

To make the computations associated with EQ. 2 easier to implement, the probability $p(\mathbf{y}|\mathbf{x}, m)$ is assumed to be independent of m so that it becomes $p(\mathbf{y}|\mathbf{x})$. This probability is then modeled as a Gaussian distribution that describes the certainty of the noise reduction process. In particular:

$$p(\mathbf{y}|\mathbf{x}) = \alpha N(\hat{\mathbf{x}}; \mathbf{x}, \sigma_{\hat{\mathbf{x}}}^2) \quad \text{EQ. 3}$$

where the distribution is defined across the cleaned feature vectors, $\hat{\mathbf{x}}$, and has a mean equal to the ideal clean feature vector, $\mathbf{x}$, and a variance, $\sigma_{\hat{\mathbf{x}}}^2$, that represents the variance associated with the noise reduction process. Note that in EQ. 3 a simplified notation has been used to represent the Gaussian distribution as a single Gaussian across the entire vector. In fact, there is a separate Gaussian, with its own mean and variance, for each dimension of the feature vector. For example:

$$p(y_l|x_l) = \alpha N(\hat{x}_l; x_l, \sigma_{\hat{x}_l}^2) \quad \text{EQ. 4}$$

-16-

where y_l is component l of the noisy feature vector, x_l is component l of the clean feature vector, $\hat{x}_l$ is component l of the cleaned feature vector, and $\sigma_{\hat{x},l}^2$ is the variance (also referred to as uncertainty) associated with the noise reduction process for component l of the feature vectors.

The prior probability, $p(\mathbf{x}|m)$, is also modeled as a collection of Gaussian distributions, one for each component of the vector, such that:

$$p(x_l|m) = N(x_l; \mu_{m,l}, \sigma_{m,l}^2) \quad \text{EQ. 5}$$

where $\mu_{m,l}$ is the mean of the prior distribution, and $\sigma_{m,l}^2$ is the variance of the distribution for component l of the feature vectors.

Combining equations 2, 4 and 5 and evaluating the integral provides:

$$p(y_l|m) = \alpha N(\hat{x}_l; \mu_{m,l}, \sigma_{m,l}^2 + \sigma_{\hat{x},l}^2) \quad \text{EQ. 6}$$

Thus, under the framework of the present invention, the variance of the distribution used in the acoustic model is a combination of the variance associated with the prior model, $\sigma_{m,l}^2$, and the variance or uncertainty associated with the noise reduction process, $\sigma_{\hat{x},l}^2$. As a result, if the uncertainty associated with the noise reduction process is high for a component, there will be little difference

-17-

between the probabilities generated for each phonetic state. As a result, the component will have little effect on the decoding of the phonetic state sequence. However, if the uncertainty of the noise
5 reduction process is small, the acoustic models for the various phonetic states will generate distinctive probabilities for the component and the component will thereby strongly influence the selection of the phonetic state.

10 Under one embodiment of the present invention, the noise reduction process utilizes a prior model of dynamic aspects of clean speech, a prior model of static aspects of clean speech, and an acoustic-distortion or acoustic-environmental model
15 that describes the relationship between clean speech, noisy speech and noise. The invention thus provides a technique for determining the uncertainty in this parametric-based noise reduction process.

In the discussion below, the noise
20 reduction method is performed in the cepstral domain. In particular, each frame of clean speech, noisy speech, and noise, is represented by a cepstral feature vector that is formed by taking the logarithm and discrete cosine transform of the frequency-domain
25 representation of the signal within the frame. Thus, in the discussion below, the method attempts to identify a noise-reduced feature vector from a noisy speech feature vector. As with Equation 3 above, a simplified notation is used below that treats each
30 vector as being represented by a single distribution.

-18-

In implementation, there is a separate distribution for each component of the vector.

To reduce the noise from the noisy vector, the minimum mean square estimate of the noise-reduced vector is determined using the conditional expectation:

$$\hat{\mathbf{x}}_t = E[\mathbf{x}_t | \mathbf{y}_t, \hat{\mathbf{x}}_{t-1}] = \int \mathbf{x}_t p(\mathbf{x}_t | \mathbf{y}_t, \hat{\mathbf{x}}_{t-1}) d\mathbf{x}_t \quad \text{EQ. 7}$$

where $\hat{\mathbf{x}}_t$ is the estimate for the noise-reduced feature vector of the current frame, $\hat{\mathbf{x}}_{t-1}$ is the estimate for the noise-reduced feature vector determined for a previous frame, $\mathbf{y}_t$ is a noisy speech feature vector for the current frame, and $\mathbf{x}_t$ is a clean speech feature vector for the current frame.

Using Bayes rule, the estimate in Equation 1 can be restated as:

$$\hat{\mathbf{x}}_t = \frac{\int \mathbf{x}_t p(\mathbf{x}_t | \hat{\mathbf{x}}_{t-1}) p(\mathbf{y}_t | \mathbf{x}_t) d\mathbf{x}_t}{p(\mathbf{y}_t)} \quad \text{EQ. 8}$$

where $p(\mathbf{x}_t | \hat{\mathbf{x}}_{t-1})$ is the conditional probability of a clean speech vector, $\mathbf{x}_t$, given the estimate of the clean speech vector at the preceding frame, $\hat{\mathbf{x}}_{t-1}$, $p(\mathbf{y}_t | \mathbf{x}_t)$ is the conditional probability of a noisy speech vector, $\mathbf{y}_t$, given a clean speech vector, $\mathbf{x}_t$, and $p(\mathbf{y}_t)$ is the probability of a noisy speech vector, $\mathbf{y}_t$.

Under one embodiment of the present invention, a mixture model is used to describe the

-19-

conditional probability $p(\mathbf{x}_t|\hat{\mathbf{x}}_{t-1})$. In terms of an equation:

$$p(\mathbf{x}_t|\hat{\mathbf{x}}_{t-1}) = \sum_{m=1}^M c_m p(\mathbf{x}_t|m, \hat{\mathbf{x}}_{t-1}) \quad \text{EQ. 9}$$

where the conditional prior probability $p(\mathbf{x}_t|m, \hat{\mathbf{x}}_{t-1})$ is
 5 the probability provided by the m th mixture component
 and c_m is a weighting factor for the m th mixture
 component.

Under one embodiment, noisy speech is
 considered to be the result of noise being added to
 10 clean speech. This provides an acoustic environment
 model otherwise known as a distortion model for the
 noisy speech. In the time domain, this addition is
 linear. However, in the cepstral feature domain,
 this addition becomes non-linear resulting in a
 15 feature vector relationship of:

$$\mathbf{y}_t = \mathbf{x}_t + g(\mathbf{n}_t - \mathbf{x}_t) + \mathbf{r} \quad \text{EQ. 10}$$

where

$$g(\mathbf{n}_t - \mathbf{x}_t) = \log(1 + e^{\mathbf{n}_t - \mathbf{x}_t}) \quad \text{EQ. 11}$$

and $\mathbf{r}$ is a residual.

20 To make the expression of Equation 10
 easier to use, the non-linearity is linearized using
 a truncated Taylor series approximation. In
 particular, a zero-th order Taylor series expansion
 on $g(\mathbf{n}_t - \mathbf{x}_t)$ at $\mathbf{x}_t = \mathbf{x}_0$ is used to form the
 25 approximation:

$$\mathbf{y}_t \approx \mathbf{x}_t + g(\bar{\mathbf{n}} - \mathbf{x}_0) + \mathbf{r} \quad \text{EQ. 12}$$

-20-

where $g(\bar{\mathbf{n}} - \mathbf{x}_0)$ is the zero-th order Taylor series expansion at expansion point $\mathbf{x}_0$, and $\bar{\mathbf{n}}$ is an estimate of the noise in the current frame. In most embodiments, the expansion point, $\mathbf{x}_0$, is selected as the estimate of the clean speech signal determined during a previous iteration of the noise-reduction process, and $\bar{\mathbf{n}}$ is determined using a recursive technique. Those skilled in the art will recognize that the recursive technique for noise estimation is one example of a number of different noise estimation techniques that may be utilized with the present invention.

Based on the relationship shown in Equation 12, the conditional probability $p(\mathbf{y}_i | \mathbf{x}_i)$ of Equation 8 becomes an acoustic environment probability $p(\mathbf{y}_i | \mathbf{x}_i, \bar{\mathbf{n}}_i)$. Using this acoustic environment probability and Equation 9 above, Equation 8 can be re-written as:

$$\hat{\mathbf{x}}_i = \frac{\sum_{m=1}^M c_m \int \mathbf{x}_i p(\mathbf{x}_i | m, \hat{\mathbf{x}}_{i-1}) p(\mathbf{y}_i | \mathbf{x}_i, \bar{\mathbf{n}}_i) d\mathbf{x}_i}{p(\mathbf{y}_i)} \quad \text{EQ. 13}$$

To compute the integral of Equation 13, the conditional prior, $p(\mathbf{x}_i | m, \hat{\mathbf{x}}_{i-1})$, is first evaluated. To make this evaluation easier, it is recognized that:

$$p(\mathbf{x}_i | m, \hat{\mathbf{x}}_{i-1}) \propto p(\mathbf{x}_i, \Delta \mathbf{x}_i | m) \quad \text{EQ. 14}$$

where

$$\Delta \mathbf{x}_i = \mathbf{x}_i - \hat{\mathbf{x}}_{i-1} \quad \text{EQ. 15}$$

-21-

Thus, in Equation 14, $\mathbf{x}_i$ represents the static aspects of clean speech and $\Delta\mathbf{x}_i$ represents the dynamic aspects of clean speech.

To further simplify the evaluation of the conditional prior, it is assumed that the static aspects of clean speech are uncorrelated to the dynamic aspects of clean speech. As a result:

$$p(\mathbf{x}_i, \Delta\mathbf{x}_i | m) = p(\mathbf{x}_i | m) p(\Delta\mathbf{x}_i | m) \quad \text{EQ. 16}$$

Under one embodiment, the prior probability for the static aspects of clean speech and the prior probability for the dynamic aspects of clean speech are each modeled as Gaussian distributions such that:

$$p(\mathbf{x}_i, \Delta\mathbf{x}_i | m) = N(\mathbf{x}_i; \boldsymbol{\mu}_m^x, \boldsymbol{\Phi}_m^x) N(\Delta\mathbf{x}_i; \boldsymbol{\mu}_m^{\Delta x}, \boldsymbol{\Phi}_m^{\Delta x}) \quad \text{EQ. 17}$$

where $\boldsymbol{\mu}_m^x$ and $\boldsymbol{\Phi}_m^x$ are the mean and variance of the prior model for the static aspects of clean speech, and $\boldsymbol{\mu}_m^{\Delta x}$ and $\boldsymbol{\Phi}_m^{\Delta x}$ are the mean and variance of the prior model for the dynamic aspects of clean speech.

Fitting the exponent in the product of the above two Gaussian distributions into the standard quadratic form in $\mathbf{x}_i$, and using Equation 14 above, produces:

$$p(\mathbf{x}_i | m, \hat{\mathbf{x}}_{i-1}) = N(\mathbf{x}_i; \boldsymbol{\mu}_m, \boldsymbol{\Phi}_m) \quad \text{EQ. 18}$$

where

$$\boldsymbol{\mu}_m = \frac{\boldsymbol{\Phi}_m^{\Delta x}}{(\boldsymbol{\Phi}_m^x + \boldsymbol{\Phi}_m^{\Delta x})} \boldsymbol{\mu}_m^x + \frac{\boldsymbol{\Phi}_m^x}{(\boldsymbol{\Phi}_m^x + \boldsymbol{\Phi}_m^{\Delta x})} (\mathbf{x}_{i-1} + \boldsymbol{\mu}_m^{\Delta x}) \quad \text{EQ. 19}$$

and

-22-

$$\Phi_m = \frac{\Phi_m^x \Phi_m^{\Delta x}}{(\Phi_m^x + \Phi_m^{\Delta x})} \quad \text{EQ. 20}$$

In many embodiments, the acoustic environment probability, $p(\mathbf{y}_t | \mathbf{x}_t, \bar{\mathbf{n}}_t)$, is also modeled as a Gaussian distribution producing an acoustic model of the form:

$$p(\mathbf{y}_t | \mathbf{x}_t, \mathbf{n}_t) = N(\mathbf{y}_t; \mathbf{x}_t + g(\bar{\mathbf{n}} - \mathbf{x}_0), \Psi) \quad \text{EQ. 21}$$

where Ψ is a fixed diagonal covariance matrix that is an approximation to the covariance matrix for the residual $\mathbf{r}$.

Thus, the integral in the numerator of Equation 13 becomes:

$$I_m = \int \mathbf{x}_t N(\mathbf{x}_t; \mu_m, \Phi_m) N(\mathbf{y}_t; \mathbf{x}_t + g(\bar{\mathbf{n}} - \mathbf{x}_0), \Psi) d\mathbf{x}_t \quad \text{EQ. 22}$$

Evaluating the integral produces:

$$I_m = [v_1(m) \mu_m^x + v_2(m)(\mathbf{x}_{t-1} + \mu_m^{\Delta x}) + w_2(m)(\mathbf{y}_t - g(\bar{\mathbf{n}} - \mathbf{x}_0))] N_m(\mathbf{y}_t) \quad \text{EQ. 23}$$

where

$$N_m(\mathbf{y}_t) = N(\mathbf{y}_t; \mu_m^x + g(\bar{\mathbf{n}} - \mathbf{x}_0), \Phi_m^x + \Psi) \quad \text{EQ. 24}$$

$$v_1(m) = \frac{\Psi}{(\Phi_m^x + \Psi)} \frac{\Phi_m^{\Delta x}}{(\Phi_m^x + \Phi_m^{\Delta x})} \quad \text{EQ. 25}$$

$$v_2(m) = \frac{\Psi}{(\Phi_m^x + \Psi)} \frac{\Phi_m^x}{(\Phi_m^x + \Phi_m^{\Delta x})} \quad \text{EQ. 26}$$

$$w_2(m) = \mathbf{I} - \frac{\Psi}{(\Phi_m^x + \Psi)} \quad \text{EQ. 27}$$

The Gaussian distribution in Equation 25 can be shown to be the likelihood of observation $\mathbf{y}_t$,

-23-

given the m -th component in the clean speech model under the zero-th order approximation made in Equation 15. That is:

$$p(\mathbf{y}_t|m) \approx N_m(\mathbf{y}_t) \quad \text{EQ. 28}$$

5 As a result, the denominator in Equation 13 can be determined as:

$$p(\mathbf{y}_t) = \sum_{m=1}^M c_m p(\mathbf{y}_t|m) \approx \sum_{m=1}^M c_m N_m(\mathbf{y}_t) \quad \text{EQ. 29}$$

Combining Equations 13, 23, and 29 produces the estimator for $\mathbf{x}_t$:

$$10 \quad \mathbf{x}_t = \sum_{m=1}^M \gamma_m(\mathbf{y}_t) \left[v_1(m) \boldsymbol{\mu}_m^x + v_2(m) (\mathbf{x}_{t-1} + \boldsymbol{\mu}_m^{\Delta x}) + w_2(m) (\mathbf{y}_t - g(\bar{\mathbf{n}} - \mathbf{x}_0)) \right] \quad \text{EQ. 30}$$

where

$$\gamma_m(\mathbf{y}_t) = \frac{c_m N_m(\mathbf{y}_t)}{\sum_{m=1}^M c_m N_m(\mathbf{y}_t)} \quad \text{EQ. 31}$$

15 Each summand in Equation 30 is a mixture component contribution to the final reduced-noise vector. Each summand is formed by multiplying the weighted sum of three terms by the relative probability of the noisy feature vector, $\mathbf{y}_t$, occurring
 20 given the mixture component. The first term is the prediction of the reduced-noise vector taken from the mean of the static aspects of the prior clean speech model, $\boldsymbol{\mu}_m^x$. The second term is the prediction for the reduced-noise vector based on the estimate of the
 25 reduced-noise vector for the preceding frame and the

-24-

mean of the dynamic aspects of the prior clean speech model, $(\mathbf{x}_{t-1} + \boldsymbol{\mu}_m^{\Delta x})$. The last term, $\mathbf{y}_t - g(\bar{\mathbf{n}} - \mathbf{x}_0)$, is the prediction for the noise-reduced clean-speech vector based on the acoustic distortion model in absence of
5 any prior information.

Thus, the estimator of Equation 30 utilizes a prior model of static aspects of clean speech, a prior model of dynamic aspects of clean speech and an acoustic distortion model. As a result, the sequence
10 of noise-reduced vectors produced using the estimator of Equation 30 contain fewer discontinuities and appear more natural.

In addition, the weighting values applied to each of the terms in each summand are such that as
15 the variance of one of the terms increases, its relative weighting decreases. Thus, as a term becomes more accurate in its description of the noise-reduced value, its weighting increases relative to the weighting of the other terms.

For example, the weighting value applied to
20 the value predicted by the prior model for static aspects of clean speech is a function of the variance of the prior model for dynamic aspects of clean speech. As a result, when the variance of the model
25 for the dynamic aspects of clean speech increases relative to the variance of the model for the static aspects of clean speech, the relative weighting of the value predicted by the prior model for static aspects increases.

-25-

A similar effect occurs for the weighting of the value predicted by the prior model for the dynamic aspects of clean speech, since this weighting value is a function of the variance of the prior
 5 model for the static aspects of clean speech.

Before the estimator of Equation 31 can be used, the parameters for the various models must be trained. Under one embodiment, an EM algorithm is used to train the mean and covariance parameters μ_m^x ,
 10 Φ_m^x , $\mu_m^{\Delta x}$ and $\Phi_m^{\Delta x}$ in the cepstral domain. During this training process, the mixture component weights, c_m , are also trained. Techniques for performing such training are well known in the art. The fixed covariance matrix, Ψ , is estimated by pooling the
 15 training data with all available signal to noise ratios.

Under one embodiment, the noise reduced vector, $\mathbf{x}_t$, is determined in an iterative fashion such that the value of the noise-reduced vector in a
 20 current iteration is used as the expansion point, $\mathbf{x}_0$, of the Taylor Series approximation, $g(\bar{\mathbf{n}} - \mathbf{x}_0)$, in the next iteration. For the first iteration, the expansion point for the Taylor Series approximation is the mixture component mean, μ_m^x , that provides the
 25 highest probability for the noisy vector using the normal distribution:

$$N(\mathbf{y}; \mu_m^x + g(\bar{\mathbf{n}} - \mathbf{x}_0), \Psi) \quad \text{EQ. 32}$$

-26-

At each iteration, the noise-reduced vector is then calculated as:

$$\hat{\mathbf{x}}_t^{j+1} = \sum_{m=1}^M \gamma_m^j(\mathbf{y}_t) \left[v_1(m) \boldsymbol{\mu}_m^x + v_2(m) (\hat{\mathbf{x}}_{t-1} + \boldsymbol{\mu}_m^{\Delta x}) + w_2(m) (\mathbf{y}_t - g(\bar{\mathbf{n}} - \hat{\mathbf{x}}_t^j)) \right]$$

EQ. 33

5 where j represents an iteration counter and

$$\gamma_m^j(\mathbf{y}_t) = \frac{c_m N(\mathbf{y}_t; \boldsymbol{\mu}_m^x + g(\bar{\mathbf{n}} - \hat{\mathbf{x}}_t^j), \boldsymbol{\Phi}_m^x + \boldsymbol{\Psi})}{\sum_{m=1}^M c_m N(\mathbf{y}_t; \boldsymbol{\mu}_m^x + g(\bar{\mathbf{n}} - \hat{\mathbf{x}}_t^j), \boldsymbol{\Phi}_m^x + \boldsymbol{\Psi})}$$

EQ. 34

Thus, the expectation $\boldsymbol{\mu}_{\hat{\mathbf{x}}_t}$ for the noise-reduced vector is:

$$\boldsymbol{\mu}_{\hat{\mathbf{x}}_t} = \hat{\mathbf{x}}_t^J$$

EQ. 35

where J is the total number of iterations.

Under the present invention, the uncertainty associated with the noise reduction process represented by the iteration of equation 33 is also calculated so that it can be used during recognition. The uncertainty is the variance $\Sigma_{\hat{\mathbf{x}}_t}$ in $\boldsymbol{\mu}_{\hat{\mathbf{x}}_t}$, which is defined as:

$$\Sigma_{\hat{\mathbf{x}}_t} = E \left[x_t^2 | y_t \right] - \boldsymbol{\mu}_{\hat{\mathbf{x}}_t}^2$$

EQ. 36

where

$$E \left[x_t^2 | y_t \right] \approx \frac{\sum_{m=1}^M c_m \int_{\mathbf{x}_t} \overbrace{\mathbf{x}_t^2}^{I_m(y_t)} p(\mathbf{x}_t | m, \hat{\mathbf{x}}_{t-1}) p(\mathbf{y}_t | \mathbf{x}_t, \bar{\mathbf{n}}_t) d\mathbf{x}_t}{p(\mathbf{y}_t)}$$

EQ. 37

-27-

Using the approximations and models of equations 14-21 above, the integral I_m , becomes:

$$\begin{aligned} I_m &= \int \mathbf{x}_t^2 N(\mathbf{x}_t; \boldsymbol{\mu}_m, \boldsymbol{\Phi}_m) N(\mathbf{y}_t; \mathbf{x}_t + g(\bar{\mathbf{n}} - \mathbf{x}_0), \boldsymbol{\Psi}) d\mathbf{x}_t \\ &= \int \mathbf{x}_t^2 N(\mathbf{x}_t; \boldsymbol{\theta}_m(t), (\boldsymbol{\Phi}_m + \boldsymbol{\Psi})^{-1} \boldsymbol{\Phi}_m \boldsymbol{\Psi}) d\mathbf{x}_t \times N_m(y_t) \text{EQ. 38} \\ &= \left[(\boldsymbol{\Phi}_m + \boldsymbol{\Psi})^{-1} \boldsymbol{\Phi}_m \boldsymbol{\Psi} + \boldsymbol{\theta}_m^2(t) \right] \times N_m(y_t) \end{aligned}$$

where

$$\boldsymbol{\mu}_m = \frac{\boldsymbol{\Phi}_m^{\Delta x}}{(\boldsymbol{\Phi}_m^x + \boldsymbol{\Phi}_m^{\Delta x})} \boldsymbol{\mu}_m^x + \frac{\boldsymbol{\Phi}_m^x}{(\boldsymbol{\Phi}_m^x + \boldsymbol{\Phi}_m^{\Delta x})} (\mathbf{x}_{t-1} + \boldsymbol{\mu}_m^{\Delta x}) \text{EQ. 39}$$

$$\boldsymbol{\Phi}_m = \frac{\boldsymbol{\Phi}_m^x \boldsymbol{\Phi}_m^{\Delta x}}{(\boldsymbol{\Phi}_m^x + \boldsymbol{\Phi}_m^{\Delta x})} \text{EQ. 40}$$

$$\boldsymbol{\theta}_m(t) = (\boldsymbol{\Phi}_m + \boldsymbol{\Psi})^{-1} [\boldsymbol{\Psi} \boldsymbol{\mu}_m + \boldsymbol{\Phi}_m (\mathbf{y}_t - g(\bar{\mathbf{n}} - \mathbf{x}_0))] \text{EQ. 41}$$

$$N_m(\mathbf{y}_t) = N(\mathbf{y}_t; \boldsymbol{\mu}_m + g(\bar{\mathbf{n}} - \mathbf{x}_0), \boldsymbol{\Phi}_m + \boldsymbol{\Psi}) \text{EQ. 42}$$

Substituting the result of equation 38 into equation 37 we obtain:

$$E[x_t^2 | y_t] \approx \sum_{m=1}^M \gamma_m(\mathbf{y}_t) \left[(\boldsymbol{\Phi}_m + \boldsymbol{\Psi})^{-1} \boldsymbol{\Phi}_m \boldsymbol{\Psi} + \boldsymbol{\theta}_m^2(t) \right] \text{EQ. 43}$$

where

$$\gamma_m(\mathbf{y}_t) = \frac{c_m N_m(\mathbf{y}_t)}{\sum_{m=1}^M c_m N_m(\mathbf{y}_t)} \text{EQ. 44}$$

The value calculated in Equation 43 is then used together with the square of the expectation for the noise reduced vector calculated in equations 33-35 to determine the uncertainty of the noise reduction process using equation 36 above. Note that in Equation 36, $\Sigma_{\hat{x}_t}$ is a vector representation of the

-28-

uncertainty and that it represents a separate uncertainty $\sigma_{\hat{x}_i}^2$ for each component of the static noise reduced vector $\mu_{\hat{x}_i}$.

Equations 33-35 above determine the static components of the noise reduced vector. Under 5 embodiments of the present invention, differentials of these static components that describe how the components change between frames can also be determined. These differentials are referred to 10 generically as dynamic features. Under the present invention, two types of dynamic features are determined. One is the delta feature that indicates changes between frames, and the other is the acceleration feature that indicates the rate of 15 change of the delta features. Under one embodiment, the delta features are calculated as:

$$\Delta \hat{\mathbf{x}}_t = \sum_{\tau=-K}^K \omega_{\tau} \hat{\mathbf{x}}_{t+\tau} \quad \text{EQ. 45}$$

where $\Delta \hat{\mathbf{x}}_t$ is the dynamic feature for the noise reduced vector, K indicates the number of neighboring frames 20 that are to be used to determine the dynamic features, and ω_{τ} are fixed weights that provide relative weighting to the neighboring frames. Note that the vector notation is used in equation 45 to indicate that a separate calculation is performed for 25 each component of the noise-reduced feature vectors.

Acceleration features can also be determined based on the delta features calculated in Equation 45. Specifically:

-29-

$$\Delta^2 \hat{\mathbf{x}}_i = \sum_{\tau=-K}^K \nu_{\tau} \Delta \hat{\mathbf{x}}_{i+\tau} \quad \text{EQ. 46}$$

where $\Delta^2 \hat{\mathbf{x}}_i$ is the acceleration feature for the noise reduced vector, K indicates the number of neighboring frames that are to be used to determine the dynamic features, and ν_{τ} are fixed weights that provide relative weighting to the neighboring frames.

The delta and acceleration features are provided to the decoder along with the static noise-reduced features. In addition, the uncertainty of the delta and acceleration features are calculated and provided to the decoder. Specifically the uncertainty of the delta and acceleration features are calculated as:

$$\Sigma_{\Delta \hat{\mathbf{x}}_i} = \sum_{\tau=-K}^K \omega_{\tau}^2 \Sigma_{\hat{\mathbf{x}}_i} \quad \text{EQ. 47}$$

$$\Sigma_{\Delta^2 \hat{\mathbf{x}}_i} = \sum_{\tau=-K}^K \nu_{\tau}^2 \Sigma_{\Delta \hat{\mathbf{x}}_i} \quad \text{EQ. 48}$$

where $\Sigma_{\hat{\mathbf{x}}_i}$ is the uncertainty calculated in Equation 36 above, $\Sigma_{\Delta \hat{\mathbf{x}}_i}$ is the uncertainty in the delta features, $\Sigma_{\Delta^2 \hat{\mathbf{x}}_i}$ is the uncertainty in the acceleration features and ω_{τ} and ν_{τ} are the same weights used in Equations 45 and 46.

The uncertainty of the static features, the delta features and the acceleration features are each used in Equation 6 to determine the probability of corresponding features of the noisy feature vector during decoding. In particular, each component of

-30-

the delta feature and acceleration feature calculated in Equations 45 and 46 is treated as a separate component of the noise-reduced vector and the corresponding uncertainty associated with each component of the delta vector and each acceleration vector is used as an uncertainty $\sigma_{x_i}^2$, in Equation 6.

A method and system for using the present invention in speech recognition is shown in the flow diagram of FIG. 3 and the block diagram of FIG. 4. The method begins in step 300 of FIG. 3, where a noisy speech signal is converted into a sequence of feature vectors. To do this, a microphone 404 of FIG. 4, converts audio waves from a speaker 400 and one or more additive noise sources 402 into electrical signals. The electrical signals are then sampled by an analog-to-digital converter 406 to generate a sequence of digital values, which are grouped into frames of values by a frame constructor 408. In one embodiment, A-to-D converter 406 samples the analog signal at 16 kHz and 16 bits per sample, thereby creating 32 kilobytes of speech data per second and frame constructor 408 creates a new frame every 10 milliseconds that includes 25 milliseconds worth of data.

Each frame of data provided by frame constructor 408 is converted into a feature vector by a feature extractor 410. Methods for identifying such feature vectors are well known in the art and

-31-

include 13-dimensional Mel-Frequency Cepstrum Coefficients (MFCC) extraction.

In step 302 of FIG. 3, the feature vectors for the noisy speech signal are provided to a noise estimation module 411 in FIG. 4. Noise estimation module 411 estimates the noise in the current frame and provides a feature vector representing the noise estimate together with the noisy speech signal to a noise reduction module 412.

In step 304 of FIG. 3, noise reduction module 412 uses Equations 33-35, 45 and 46 above and the model parameters of those equations, which are stored in noise reduction parameter storage 413, to produce a sequence of noise-reduced feature vectors from the sequence of noisy feature vectors. In particular, noise reduction module 412 uses an acoustic distortion model as described above. In addition, noise reduction module 412 uses equations 36-44, 47 and 48 to determine an uncertainty associated with forming the noise reduced feature vectors.

The output of noise reduction module 412 is a series of noise-reduced feature vectors and a corresponding series of uncertainties associated with those noise-reduced feature vectors. If the input signal is a training signal, the series of noise-reduced feature vectors are provided to a trainer 424, which uses the noise-reduced feature vectors and a training text 426 to train an acoustic model 418. Techniques for training such models are known in the

-32-

art and a description of them is not required for an understanding of the present invention.

If the input signal is a test signal, the noise-reduced feature vectors and the corresponding
5 uncertainties are provided to a decoder 414, which identifies a most likely sequence of words based on the stream of feature vectors, the stream of uncertainties, a lexicon 415, a language model 416, and the acoustic model 418 as shown by step 306 of
10 FIG. 3. Specifically, the decoding utilizes equation 6 above.

The most probable sequence of hypothesis words is provided to a confidence measure module 420. Confidence measure module 420 identifies which words
15 are most likely to have been improperly identified by the speech recognizer, based in part on a secondary acoustic model(not shown). Confidence measure module 420 then provides the sequence of hypothesis words to an output module 422 along with identifiers
20 indicating which words may have been improperly identified. Those skilled in the art will recognize that confidence measure module 420 is not necessary for the practice of the present invention.

Although FIG. 4 depicts a speech
25 recognition system, the present invention may be used in any pattern recognition system and is not limited to speech.

Although the present invention has been described with reference to particular embodiments,
30 workers skilled in the art will recognize that

-33-

changes may be made in form and detail without departing from the spirit and scope of the invention.

-34-

WHAT IS CLAIMED IS:

1. A method of reducing noise in a noisy signal, the method comprising:
reducing noise in a representation of a portion of the noisy signal to produce a representation of a portion of a noise-reduced signal using an acoustic-environment model; and
identifying an uncertainty associated with reducing the noise.
2. The method of claim 1 further comprising
using the uncertainty to adjust a probability distribution to form a modified probability distribution; and
using the representation of a portion of the noise-reduced signal and the modified probability distribution to decode a pattern state.
3. The method of claim 1 wherein the representation of a portion of a noise-reduced signal comprises a representation of delta features of the noise-reduced signal.
4. The method of claim 3 wherein identifying an uncertainty comprises identifying an uncertainty associated with a delta feature of the noise-reduced signal.

-35-

5. The method of claim 1 wherein the representation of a portion of a noise-reduced signal comprises a representation of acceleration features of the noise-reduced signal.

6. The method of claim 5 wherein identifying an uncertainty comprises identifying an uncertainty associated with an acceleration feature of the noise-reduced signal.

7. The method of claim 1 wherein reducing noise comprises using a prior model of dynamic aspects of clean signals.

8. A computer-readable medium having computer-executable instructions for performing steps comprising:

converting a frame of a noisy signal into a feature vector comprising at least two components;

reducing noise in a component of the feature vector for the noisy signal to produce a component of a feature vector for a cleaned signal by utilizing an acoustic distortion model; and

identifying an uncertainty associated with reducing the noise from the component.

-36-

9. The computer-readable medium of claim 8 wherein the computer-executable instructions perform further steps comprising:

determining a probability component of a probability for a phonetic state based in part on the component for the cleaned signal and the uncertainty associated with reducing the noise from the component; and

using the probability component to determine the probability of the phonetic state regardless of the value of the uncertainty.

10. The computer-readable medium of claim 9 wherein determining a probability component comprises defining a probability distribution based in part on the uncertainty.

11. The computer-readable medium of claim 10 wherein defining a probability distribution comprises adding the uncertainty to a variance of a probability distribution.

12. The computer-readable medium of claim 8 wherein the computer-executable instructions perform a further step comprising determining a delta feature of the cleaned signal.

-37-

13. The computer-readable medium of claim 12 wherein the computer-executable instructions perform a further step comprising determining an uncertainty of the delta feature of the cleaned signal.

14. The computer-readable medium of claim 8 wherein the computer-executable instructions perform a further step comprising determining an acceleration feature of the cleaned signal.

15. The computer-readable medium of claim 14 wherein the computer-executable instructions perform a further step comprising determining an uncertainty of the acceleration feature of the cleaned signal.

16. The computer-readable medium of claim 8 wherein reducing noise from a component further comprises utilizing a prior model of dynamic aspects of a clean signal.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 03/16032

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G10L15/20 G10L15/02 G10L21/02

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G10L

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

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

INSPEC, EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	DROPPO J ET AL: "Uncertainty decoding with SPLICE for noise robust speech recognition" PROCEEDINGS OF INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH AND SIGNAL PROCESSING (CAT. NO.02CH37334) (ICASSP'02), ORLANDO, FL, USA, 13-17 MAY 2002, pages I-57-60 vol.1, XP002254998 Piscataway, NJ, USA, IEEE, USA ISBN: 0-7803-7402-9 the whole document --- -/--	1-8, 12-16

☒ Further documents are listed in the continuation of box C.☐ Patent family members are listed in annex.

* Special categories of cited documents:

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

E earlier document but published on or after the international filing date

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

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

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

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

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

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

& document member of the same patent family

Date of the actual completion of the international search

19 September 2003

Date of mailing of the international search report

01/10/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Santos Luque, R

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 03/16032

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	LI DENG ET AL: "A Bayesian approach to speech feature enhancement using the dynamic cepstral prior" PROCEEDINGS OF INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH AND SIGNAL PROCESSING (CAT. NO.02CH37334) (ICASSP'02), ORLANDO, FL, USA, 13-17 MAY 2002, pages I-829-32 vol.1, XP002254999 Piscataway, NJ, USA, IEEE, USA ISBN: 0-7803-7402-9 the whole document ---	1-16
A	JASHA DROPPA ET AL: "Evaluation of the SPLICE Algorithm on the Aurora2 Database" 7TH EUROPEAN CONFERENCE ON SPEECH COMMUNICATION AND TECHNOLOGY (EUROSPEECH 2001). PROCEEDINGS OF EUROSPEECH 2001, AALBORG, DENMARK, vol. 1, 3 - 7 September 2001, pages 217-220, XP007004815 abstract -----	1-16