

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 November 2008 (13.11.2008)

PCT

(10) International Publication Number
WO 2008/137616 A1

(51) International Patent Classification:
G10L 15/06 (2006.01)

(74) Agents: SUNSTEIN, Bruce, D. et al.; Bromberg & Sunstein, LLP, 125 Summer Street, Boston, MA 02110 (US).

(21) International Application Number:
PCT/US2008/062322

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date: 2 May 2008 (02.05.2008)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/915,971 4 May 2007 (04.05.2007) US

(71) Applicant (for all designated States except US): NUANCE COMMUNICATIONS, INC. [US/US]; 1 Wayside Road, Suite 100, Burlington, MA 01803-4613 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): DESHMUKH, Neeraj [IN/US]; 4414 Stearns Hill Road, Waltham, MA 02451 (US). ZHAN, Puming [CN/US]; 5 Fischer Path, Acton, MD 01720 (US).

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

Published:

— with international search report

(54) Title: MULTI-CLASS CONSTRAINED MAXIMUM LIKELIHOOD LINEAR REGRESSION




FIG. 2

(57) Abstract: A method of speech recognition converts an unknown speech input into a stream of representative features. The feature stream is transformed based on speaker dependent adaptation of multi-class feature models. Then automatic speech recognition is used to compare the transformed feature stream to multi-class speaker independent acoustic models to generate an output representative of the unknown speech input.

Multi-Class Constrained Maximum Likelihood Linear Regression

[0001] This application claims priority from U.S. Provisional Patent Application 60/915,971, filed May 4, 2007, which is incorporated herein by reference.

Field of the Invention

[0002] The invention generally relates to automatic speech recognition (ASR), and more specifically, to adaptation of feature space transformations for ASR.

Background Art

[0003] A speech recognition system determines representative text corresponding to input speech. Typically, the input speech is processed into a sequence of digital frames. Each frame can be thought of as a multi-dimensional vector that represents various characteristics of the speech signal present during a short time window of the speech. In a continuous recognition system, variable numbers of frames are organized as “utterances” representing a period of speech followed by a pause, which in real life loosely corresponds to a spoken sentence or phrase.

[0004] The system compares the input utterances to find acoustic models that best match the frame characteristics and determine corresponding representative text associated with the acoustic models. Modern acoustic models typically use Hidden Markov Models and they model speech sounds (usually phonemes) using mixtures of Gaussians. Often these phoneme models represent phonemes in specific contexts, referred to as PELs (Phonetic Elements), e.g. triphones or phonemes with known left and/or right contexts.

[0005] State sequence models can be scaled up to represent words as connected sequences of acoustically modeled phonemes, and phrases or sentences as connected sequences of words. When the models are organized together as words, phrases, and sentences, additional language-related information is also typically incorporated into the models in the form of language modeling.

[0006] The words or phrases associated with the best matching model structures are referred to as recognition candidates or hypotheses. A system may produce a single best recognition candidate – the recognition result – or a list of several hypotheses, referred to as an N-best list. Further details regarding continuous speech recognition are provided in U.S. Patent No. 5,794,189, entitled “Continuous Speech Recognition,” and U.S. Patent No. 6,167,377, entitled “Speech Recognition Language Models,” the contents of which are incorporated herein by reference.

[0007] Speech recognition can be classified as being either speaker independent or speaker dependent. The models in a speaker dependent system are specific to an individual user. Known speech inputs from the user are used to adapt a set of initially generic recognition models to specific speech characteristics of that user. The speaker adapted models form the basis for a user profile to perform speaker dependent or speaker adapted speech recognition for that user.

[0008] Speaker dependent systems traditionally use an enrollment procedure to initially create a user profile and a corresponding set of adapted models before a new user can use the system to recognize unknown inputs. During the enrollment procedure, the new user inputs speech corresponding to a known source script that is provided. During this enrollment process, the acoustic models are adapted to the specific speech characteristics of that user. These adapted models form the main portion of the user profile and are used to perform post-enrollment speech recognition for that user. Further details regarding speech recognition enrollment are provided in U.S. Patent No. 6,424,943, entitled “Non-Interactive Enrollment in Speech Recognition,” which is incorporated herein by reference.

[0009] Speaker dependent speech recognition systems running on modern desktops use adaptation at many simultaneous levels to improve accuracy and recognition speed. Some of these techniques, such as cepstral normalization, histogram normalization, or speaker adaptive training (SAT) operate directly on the input speech feature stream. Others, such as maximum likelihood linear regression (MLLR) and maximum a posteriori parameter estimation (MAP) operate by transforming speech recognition models to better fit the incoming signal.

[0010] One typical use of MLLR adaptation in speech recognition has been to transform sets of Gaussian mixtures models which share some property, such as the same center phoneme. These sets are referred to as “classes,” and MLLR in this context can be thought of as class-based MLLR.

[0011] One specific form of MLLR, constrained MLLR (cMLLR) has been used for several years in state of the art recognition systems. In contrast to generic MLLR, cMLLR constrains the linear transformations to modify both the model means and variances in a consistent fashion. The resulting transformation can then be inverted and applied in the feature space rather than in the model space. In specific example of the use of cMLLR uses an online unsupervised feature space adaptation (OUFA) technique, described further in U.S. Patent Application 11/478,837, entitled “Non-Enrolled Continuous Dictation,” the contents of which are incorporated herein by reference.

[0012] Most server-based speech recognition systems avoid most types of adaptation, and particularly model-space adaptation, but Figure 1 shows an example of a multiple speaker application using OUFA and a single class cMLLR transform. During initial enrollment, an enrollment speech input from a new user is used to perform MLLR adaptation of an initial speaker independent acoustic model (SIAM) to generate a speaker dependent acoustic model (SDAM) so that the system has one such SDAM for each of the N speakers registered with the system, $SDAM_1$ - $SDAM_N$. The enrollment speech input is also used to perform single class cMLLR adaptation wherein the inverse cMLLR transform is used to adapt feature processing blocks—speaker dependent SAT models (one of the blocks SAT_1 - SAT_N) and OUFA models (one of the blocks $OUFA_1$ - $OUFA_N$) for that speaker.

[0013] After enrollment, unknown input speech from a given user is initially processed by a speaker dependent front end (one of the blocks $SDFE_1$ - $SDFE_N$) to produce a set of speech features representative of the speech input. These features are then processed by the speaker dependent blocks SAT and OUFA for that user, and input to the recognition engine. The recognition engine compares the input speech features to the SDAM for that

user as constrained by the recognition language model and search algorithm to produce a representative recognition output.

Summary of the Invention

[0014] A method of speech recognition converts an unknown speech input into a stream of representative features. The feature stream is transformed based on speaker dependent adaptation of multi-class feature models, for example, by constrained maximum likelihood linear regression (cMLLR). Then automatic speech recognition is used to compare the transformed feature stream to multi-class speaker independent acoustic models to generate an output representative of the unknown speech input.

[0015] In specific embodiments, the converting an unknown speech input is based on a speaker dependent front end. Transforming the feature stream may include using speaker adaptive training (SAT) and/or online unsupervised feature space adaptation (OUFA). Converting an unknown speech input and transforming the feature stream may be performed by a local client, and the automatic speech recognition may be performed by a remote server. The multi-class speaker independent acoustic models may specifically include hidden Markov models which use Gaussian mixture components. The speech recognition may specifically use scaled integer arithmetic.

Brief Description of the Drawings

[0016] Figure 1 shows an example of a multiple speaker application using OUFA and a single class cMLLR transform.

[0017] Figure 2 shows an embodiment of the present invention using multi-class cMLLR transforms.

Detailed Description of Specific Embodiments

[0018] Various embodiments of the present invention are directed to techniques to bring the performance benefits of MLLR to server-based speech recognition by computing different cMLLR transforms for different classes of PELs. The inverse transforms are applied to the input feature stream, producing as output multiple transformed copies of the

input. These transformed features are compared with (untransformed) Gaussian mixture models for each phoneme class.

[0019] Using these untransformed Gaussian mixture models allows the acoustic models to be shared among different users. This sharing substantially decreases the amount of memory needed for a speaker dependent system, both on disk and in memory. For example, in a typical implementation, the size of the transformations is 216 kB, while the size of the Gaussian mixture models is 15 Mb. A particular advantage of multi-class cMLLR is that it allows a simultaneous multi-user system to share memory for Gaussian mixture models across users. Since Gaussian mixture models are a large part of a user's memory budget, this can result in substantial memory savings in systems which recognize speech for multiple simultaneous users.

[0020] Figure 2 shows an embodiment of the present invention using multi-class cMLLR transforms. The initial enrollment is used to perform a multi-class inverse cMLLR adaptation of speaker dependent feature space processing blocks—speaker dependent SAT and OUFA modules for that speaker. The Gaussian mixture components in the Speaker Independent Acoustic Model (SIAM) are not adapted for individual speakers.

[0021] After enrollment, unknown input speech from a given user is initially processed by a speaker dependent front end (SDFE), speaker adaptive training (SAT), and online unsupervised feature space adaptation (OUFA) for that user. The resulting stream of representative speech features is input to the recognition engine. The recognition engine compares the input speech features to the system SIAM to perform speech recognition as constrained by the recognition language model and search algorithm and produce a representative recognition output.

[0022] A formal mathematical development of single class cMLLR is presented in Gales, "Maximum Likelihood Linear Transformations for HMM-Based Speech Recognition", Tech. Report, CUED/FINFENG/TR291, Cambridge Univ., 1997, which is incorporated herein by reference. From Gales, a set of equations can be developed for multi-class

cMLLR adaptation data $G_c^{(i)}$ and $k^{(i)}$:

$$G_c^{(i)} = \sum_{m=1}^{M_c} \frac{1}{\sigma_i^{(m)^2}} \sum_{\tau=1}^{T_c} \gamma_m(\tau) \zeta(\tau) \zeta(\tau)^T$$

and,

$$k^{(i)} = \sum_{m=1}^{M_c} \frac{1}{\sigma_i^{(m)^2}} \mu_i^{(m)} \sum_{\tau=1}^{T_c} \gamma_m(\tau) \zeta(\tau)^T$$

where: i is the feature dimension index

m is the Gaussian component index

$\mu_i^{(m)}$ is the i^{th} mean parameter of Gaussian m

$\sigma_i^{(m)}$ is the i^{th} variance parameter of Gaussian m

$\zeta(\tau)$ is the feature vector at time τ , and

$\gamma_m(\tau)$ is the posterior probability at time τ for Gaussian m .

[0023] One implementation consideration in multi-class cMLLR is how to handle smoothing. Smoothing is used when the training data does not contain enough examples of some classes to well specify the transformation for that class. One approach is relevance smoothing where the adaptation data $G_c^{(i)}$ and $k^{(i)}$ for each class initially is added together:

$$\bar{G}^{(i)} = \sum_{c=1}^C G_c^{(i)} \text{ and } \bar{k}^{(i)} = \sum_{c=1}^C k_c^{(i)}$$

where C = the number of classes. Then, the total statistics are used to weight the per-class adaptation data $G_c^{(i)}$ and $k^{(i)}$:

$$\tilde{G}_c^{(i)} = \frac{r}{N} \bar{G}^{(i)} + G_c^{(i)}$$

$$\tilde{k}_c^{(i)} = \frac{r}{N} \bar{k}^{(i)} + k_c^{(i)}$$

where $0 \leq r < C$, N is the total number of frame counts, and r is a relevance factor.

[0024] In some embodiments, there could be some performance losses due to replacing MLLR transforms with cMLLR. For example, some tests have shown a relatively small accuracy loss of about 1.8% relative word error rate (WER) on one specific short

messaging service (SMS) data set.

[0025] Embodiments of the invention may be implemented in any conventional computer programming language. For example, preferred embodiments may be implemented in a procedural programming language (*e.g.*, “C”) or an object oriented programming language (*e.g.*, “C++”, Python). Alternative embodiments of the invention may be implemented as pre-programmed hardware elements, other related components, or as a combination of hardware and software components.

[0026] Embodiments can be implemented as a computer program product for use with a computer system, for example, as a series of computer instructions fixed either on a tangible medium, such as a computer readable medium (*e.g.*, a diskette, CD-ROM, ROM, or fixed disk) or transmittable to a computer system, via a modem or other interface device, such as a communications adapter connected to a network over a medium. The medium may be either a tangible medium (*e.g.*, optical or analog communications lines) or a medium implemented with wireless techniques (*e.g.*, microwave, infrared or other transmission techniques). The series of computer instructions embodies all or part of the functionality previously described herein with respect to the system. Those skilled in the art should appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Furthermore, such instructions may be stored in any memory device, such as semiconductor, magnetic, optical or other memory devices, and may be transmitted using any communications technology, such as optical, infrared, microwave, or other transmission technologies. It is expected that such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation (*e.g.*, shrink wrapped software), preloaded with a computer system (*e.g.*, on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the network (*e.g.*, the Internet). Of course, some embodiments of the invention may be implemented as a combination of both software (*e.g.*, a computer program product) and hardware. Still other embodiments of the invention are implemented as entirely hardware, or entirely software (*e.g.*, a computer program product).

[0027] Although various exemplary embodiments of the invention have been disclosed, it should be apparent to those skilled in the art that various changes and modifications can be made which will achieve some of the advantages of the invention without departing from the true scope of the invention.

What is claimed is:

1. A method of speech recognition comprising:
converting an unknown speech input into a stream of representative features;
transforming the feature stream based on speaker dependent adaptation of multi-class feature models; and
using automatic speech recognition to compare the transformed feature stream to multi-class speaker independent acoustic models to generate an output representative of the unknown speech input.
2. A method according to claim 1, wherein the adaptation is based on constrained maximum likelihood linear regression (cMLLR).
3. A method according to claim 2, wherein the adaptation based on cMLLR includes use of relevance smoothing.
4. A method according to claim 1, wherein the converting an unknown speech input is based on a speaker dependent front end.
5. A method according to claim 1, wherein the transforming the feature stream includes using speaker adaptive training (SAT).
6. A method according to claim 1, wherein the transforming the feature stream includes using online unsupervised feature space adaptation (OUFA).
7. A method according to claim 1, wherein the converting an unknown speech input and the transforming the feature stream are performed by a local client, and the using automatic speech recognition is performed by a remote server.
8. A method according to claim 1, wherein the multi-class independent acoustic models include hidden Markov models.
9. A method according to claim 8, wherein the hidden Markov models use Gaussian

mixture components.

10. A method according to claim 1, wherein the speech recognition uses scaled integer arithmetic.

11. A system for speech recognition comprising:

- means for converting an unknown speech input into a stream of representative features;
- means for transforming the feature stream based on speaker dependent adaptation of multi-class feature models; and
- means for using automatic speech recognition to compare the transformed feature stream to multi-class speaker independent acoustic models to generate an output representative of the unknown speech input.

12. A system according to claim 11, wherein the adaptation in the means for transforming is based on constrained maximum likelihood linear regression (cMLLR).

13. A system according to claim 12, wherein the adaptation based on cMLLR includes use of relevance smoothing.

14. A system according to claim 11, wherein the means for converting an unknown speech input is based on a speaker dependent front end.

15. A system according to claim 11, wherein the means for transforming the feature stream includes means for using speaker adaptive training (SAT).

16. A system according to claim 11, wherein the means for transforming the feature stream includes means for using online unsupervised feature space adaptation (OUFA).

17. A system according to claim 11, wherein the means for converting an unknown speech input and the means for transforming the feature stream are at a local client, and the means for using automatic speech recognition is at a remote server.

18. A system according to claim 11, wherein the multi-class independent acoustic models include hidden Markov models.

19. A system according to claim 18, wherein the hidden Markov models use Gaussian mixture components.

20. A system according to claim 11, wherein the means for using speech recognition includes means for using scaled integer arithmetic.




FIG. 1




FIG. 2

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2008/062322

A. CLASSIFICATION OF SUBJECT MATTER
INV. G10L15/06

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)
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)

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.
Y	YONGXIN LI ET AL: "INCREMENTAL ON-LINE FEATURE SPACE MLLR ADAPTATION FOR TELEPHONY SPEECH RECOGNITION" ICSLP 2002 : 7TH INTERNATIONAL CONFERENCE ON SPOKEN LANGUAGE PROCESSING, 16 September 2002 (2002-09-16), - 20 September 2002 (2002-09-20) page 1417, XP007011703 DENVER, COLORADO ISBN: 978-1-876346-40-9 the whole document ----- -/--	1-20

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

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

10 July 2008

Date of mailing of the international search report

22/07/2008

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, Rocio

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2008/062322

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	SAM-JOO DOH ET AL: "Inter-class MLLR for speaker adaptation" ACOUSTICS, SPEECH, AND SIGNAL PROCESSING, 2000. ICASSP '00. PROCEEDING S. 2000 IEEE INTERNATIONAL CONFERENCE ON 5-9 JUNE 2000, PISCATAWAY, NJ, USA, IEEE, vol. 3, 5 June 2000 (2000-06-05), pages 1543-1546, XP010507646 ISBN: 978-0-7803-6293-2 page 1543	1-20
A	----- US 2002/091521 A1 (YUK ET AL) 11 July 2002 (2002-07-11) page 1, paragraph 2-5; figure 1	1,11
A	----- US 5 754 681 A (WATANABE ET AL) 19 May 1998 (1998-05-19) abstract	1,11
A	----- SAON G ET AL: "Linear feature space projections for speaker adaptation" 2001 IEEE INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING. PROCEEDINGS. (ICASSP), vol. 1, 7 May 2001 (2001-05-07), - 11 May 2001 (2001-05-11) pages 325-328, XP010803061 SALT LAKE CITY, UT ISBN: 978-0-7803-7041-8 page 325 page 326, left-hand column, section 2 left-hand column	1-20
A	----- GALES M J F: "Maximum likelihood linear transformations for HMM-based speech recognition" COMPUTER SPEECH AND LANGUAGE, ELSEVIER, LONDON, GB, vol. 12, no. 2, 1 April 1998 (1998-04-01), pages 75-98, XP004418764 ISSN: 0885-2308 cited in the application the whole document	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2008/062322

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US 2002091521	A1	11-07-2002	US 2006009972	A1	12-01-2006
US 5754681	A	19-05-1998	JP 2690027	B2	10-12-1997
			JP 8106295	A	23-04-1996