EXHIBIT A

In consideration of my employment at SRI International, I agree:

1. To perform the duties assigned to me to the best of my ability, and to abide faithfully by SRI policies and practices.

2. To treat as confidential all results, intermediate and terminal, of SRI research activity in which I may participate or of which I may obtain knowledge during my employment, together with all formulae, specifications, secret processes, trade secrets, and such other confidential information belonging to SRI or its clients as may come to my knowledge in the course of or incidental to my employment, and that I shall at all times recognize and protect such property rights of SRI and its clients and not disclose same to unauthorized persons. Because much of the work done by SRI for the Government is classified, I am aware that my continued employment may depend on my ability to qualify for and to maintain an appropriate Government clearance. I also agree that I will not divulge to any unauthorized persons any classified information revealed to me during the period of my employment, and that all classified material received or generated by me will be handled in accordance with SRI Security Guide. I further warrant that to the best of my knowledge *I do not at the time of my employment have in my possession, or under my control, any material which contains* "CLASSIFIED INFORMATION" as defined in U.S. Government Industrial Security directives.

3. To promptly disclose to SRI all discoveries, improvements, and inventions, including software, conceived or made by me during the period of my employment, and I agree to execute such documents, disclose and deliver all information and data, and to do all things which may be necessary or in the opinion of SRI reasonably desirable, in order to effect transfer of ownership in or to impart a full understanding of such discoveries, improvements and inventions to SRI or to its nominee and to no other. I agree to comply with every reasonable request of SRI, its nominee, or the representative of either, for assistance in obtaining and enforcing patents. I *understand that termination of this employment shall not release me from my obligations hereunder (as well as paragraph 2 above)* provided, however, that time actually spent by me in discharging these obligations after termination of my employment shall be paid for by SRI at a reasonable rate. It is, of course, understood and agreed that I accept no responsibility for any out-of-pocket fees, costs, or expenses incurred or involved in the preparation, filing or prosecution of any application for patent or in the prosecution or defense of any litigation involving the same, and that I shall be reimbursed by SRI for any expense to which I may be put at the request of it or its nominee hereunder. This agreement does not apply to an invention which fully qualifies for the exclusion under Section 2870 of the California Labor Code which is reprinted on the reverse side of this agreement. However, all such inventions must be disclosed so that a determination can be made that they do in fact qualify for exclusion. All such disclosures will be treated as confidential.

4. That with respect to the subject matter thereof, this agreement covers my entire agreement with SRI, superseding any previous oral or written understandings or agreements with SRI or any representative thereof.

5. That my employment is not for any particular term and therefore this agreement is terminable, with immediate effect, at the will of either party.

Executed at Menlo Park, California this 8 day of April, 1996

Jennie Collins
Witness to Signature

Konig
Staff Member

Print Name: YOCHAI KONIG

By: Olivia Ornat
Human Resources, for SRI International

CALIFORNIA STATE PATENT LAW

Article 3.5
Inventions Made by an Employee

§2870. Any provision in an employment agreement which provides that an employee shall assign or offer to assign any of his or her rights in an invention to his or her employer shall not apply to an invention for which no equipment, supplies, facility, or trade secret information of the employer was used and which was developed entirely on the employee's own time, and (a) which does not relate (1) to the business of the employer or (2) to the employer's actual or demonstrably anticipated research or development, or (b) which does not result from any work performed by the employee for the employer. Any provision which purports to apply to such an invention is to that extent against the public policy of this state and is to that extent void and unenforceable.

§2871. No employer shall require a provision made void and unenforceable by Section 2870 as a condition of employment or continued employment. Nothing in this article shall be construed to forbid or restrict the right of an employer to provide in contracts of employment for disclosure, provided that any such disclosures be received in confidence, of all of the employee's inventions made solely or jointly with others during the term of his or her employment, a review process by the employer to determine such issues as may arise, and for full title to certain patents and inventions to be in the United States, as required by contracts between the employer and the United States or any of its agencies.

§2872. If an employment agreement entered into after January 1, 1980, contains a provision requiring the employee to assign or offer to assign any of his or her rights in any invention to his or her employer, the employer must also, at the time the agreement is made, provide a written notification to the employee that the agreement does not apply to an invention which qualifies fully under the provisions of Section 2870. In any suit or action arising thereunder, the burden of proof shall be on the employee claiming the benefits of its provisions.

IF YOU HAVE ANY QUESTIONS ON THE ABOVE, PLEASE BE SURE THEY ARE ANSWERED BEFORE COMMENCING EMPLOYMENT.

EXHIBIT B

MEMO

TO: Yochai Konig DATE: 7-31-99

FROM: Office of the General Counsel LOC: Menlo Park

SUBJECT: Your Termination and Sensitive Information

This memo is to remind you of your obligation to hold in confidence sensitive business and technical trade secret information of SRI International which you have been exposed to in the course of your employment at SRI.

During the time that you have been with SRI, you have necessarily received information which is useful and valuable to SRI and its clients, and is not generally known to persons outside SRI. It is particularly important to SRI that this information be appropriately protected. Consequently, as stated in your Employee Agreement, subsequent to your departure from SRI you have a continuing obligation not to use or to disclose such information to anyone.

Most likely you are already aware of the specific technical and business information that SRI considers to be a trade secret. Such information may be included in technical, scientific or business records, lab notebooks, notes, reports, blueprints, drawings, software and computer programs, client and vendor lists which should be left at SRI. If you have specific questions concerning what SRI considers to be trade secret information feel free to contact your direct supervisor or the Office of the General Counsel.

If in the future you desire to use or disclose any technical or business trade secret information that may be a trade secret of SRI, please contact us for written permission to use or to disclose it.

I have received and read a copy of this letter.

Konig (Y)

Date: Aug 5, 1999

SRI Form:
NDA Separation-Nov.93

EXHIBIT C

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT D

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT E

IN THE UNITED STATES DISTRICT COURT

IN AND FOR THE DISTRICT OF DELAWARE

- - -

PERSONALIZED USER MODEL, L.L.P.,

Plaintiff,

v.

GOOGLE, INC.,

Defendant.

: CIVIL ACTION

: NO. 09-525-LPS

- - -

Wilmington, Delaware
Tuesday, January 11, 2011
Claim Construction Hearing

- - -

BEFORE: HONORABLE **LEONARD P. STARK**, U.S.D.C.J.

- - -

APPEARANCES:

MORRIS NICHOLS ARSHT & TUNNELL, LLP
BY: KAREN JACOBS LOUDEN, ESQ., and
JEREMY A. TIGAN, ESQ.

and

SNR DENTON, LLP
BY: MARK C. NELSON, ESQ.
(Dallas, Texas)

and

SNR DENTON, LLP
BY: MARC S. FRIEDMAN, ESQ.
(New York, New York)

and

Brian P. Gaffigan
Registered Merit Reporter

APPEARANCES: (Continued)

SNR DENTON, LLP
BY: JENNIFER D. BENNETT, ESQ.
(Palo Alto, California)

Counsel for Plaintiff

POTTER ANDERSON & CORROON, LLP
BY: RICHARD L. HORWITZ, ESQ.

and

QUINN EMANUEL URQUHART OLIVER & HEDGES, LLP
BY: DAVID A. PERLSON, ESQ.
(San Francisco, California)

and

QUINN EMANUEL URQUHART OLIVER & HEDGES, LLP
BY: ANDREA PALLIOS ROBERTS, ESQ.
(Redwood Shores, California)

Counsel for Defendant

- oOo -

PROCEEDINGS

(REPORTER'S NOTE: The following claim construction hearing was held in open court, beginning at 10:08 a.m.)

09:54:36 THE COURT: Good morning, everyone. Let's start
10:08:01 by putting your appearances on the record, please.
10:08:06 MS. JACOBS LOUDEN: Good morning, your Honor.
10:08:07 THE COURT: Good morning.
10:08:08 MS. JACOBS LOUDEN: For the plaintiffs, Karen
10:08:09 Jacobs Louden and Jeremy Tigan from Morris Nichols Arhst &
10:08:14 Tunnell; and I have with me here today Mark Nelson, Jennifer
10:08:17 Bennett and Marc Friedman from the firm of SNR Denton; and
10:08:21 we also have here with us today Yochai Konig who is one of
10:08:24 the inventors and a representative of the plaintiff
10:08:32 Personalized User Model.
10:08:34 Thank you, your Honor.
10:08:34 THE COURT: Thank you.
10:08:37 MR. HORWITZ: Good morning, Your Honor.
10:08:38 THE COURT: Good morning.
10:08:38 MR. HORWITZ: Rich Horwitz from Potter Anderson
10:08:47 here today for Google; and with me from Quinn Emanuel are
10:08:48 David Perlson and Andrea Roberts at counsel table; and then
10:08:50 behind the table from Google, in-house counsel, Laura
10:08:53 Majerus and John LaBarre.
10:08:58 THE COURT: Well, welcome to all of you. So
10:09:06 we're here this morning for the Markman hearing. We
10:09:13 assigned both sides 90 minutes. Have you any suggestions
10:09:17 as to how we split that time up and actually proceed?
10:09:22 Mr. Nelson.

10:09:32 MR. NELSON: Yes, your Honor. The plaintiffs
10:09:35 and defendants -- or plaintiff and defendant had talked
10:09:37 earlier about proceeding, and I guess if you guys were still
10:09:40 in agreement, we were just going to proceed with our
10:09:43 affirmative presentation, reserving some time for rebuttal,
10:09:46 and then they were going to proceed with their presentation.
10:09:48 THE COURT: Okay. That's fine.
10:09:56 MR. PERLSON: That's fine.
10:09:58 THE COURT: Okay. Then we'll hear first from
10:09:59 the plaintiff.
10:09:59 MR. NELSON: Before we start, may we distribute
10:09:59 some materials?
10:09:59 THE COURT: That would be helpful, yes. Thank
10:09:44 you.
10:09:44 MS. JACOBS LOUDEN: Thank you, your Honor.
10:09:47 (Binders passed forward.)
10:10:08 MR. NELSON: Good morning, your Honor.
10:10:09 THE COURT: Good morning.
10:10:09 MR. NELSON: We're here today to talk about the
10:10:12 claim construction of Personalized User Model, LLP or PUM,
10:10:16 or "Pum" for short, versus Google.
10:10:21 For the Court's reference, I have the claims,
10:10:23 the three dependent claims at issue just on the board there.
10:10:27 So if the Court wants to look at the claims in total, the
10:10:38 Court may do so. And I may refer to those every once in

10:10:38 awhile.
10:10:54 We have a lot of slides, as you have now seen.
10:10:58 What I plan to do is to try to hit the real highlights.
10:10:42 I'll probably skip through a lot of the slides. There is
10:10:44 some legal stuff in there that is in the briefs. There are
10:10:47 some other things in there that are sort of text summaries
10:10:50 of the argument, so I'll try to work through it rather
10:10:53 quickly, skipping some of the slides.
10:10:55 The presentation is organized into basically
10:10:57 seven parts. We've divided the accused groups of terms in
10:11:03 order of steps and antecedent basis into sort of seven
10:11:07 somewhat related areas, and so the presentation is organized
10:11:10 into those seven areas and individual disputes about terms
10:11:14 or phrases within those seven areas are contained with that
10:11:18 part, so that is how it is organized.
10:11:20 Additionally, it's organized as Phillips
10:11:23 teaches. First, we look at the claims. Then we look at the
10:11:24 specification. Then we look at prosecution history, if it's
10:11:31 available or relevant. Then we look at extrinsic evidence,
10:11:39 to the extent that it's dictionaries or treatises at or near
10:11:39 the time of the patent. And then, finally, inventor
10:11:49 testimony, to the extent it's relevant. We don't think it
10:11:46 is, as we said in the brief. We don't have any I don't
10:11:49 think in our affirmative presentation but for one slide that
10:11:59 doesn't relate to the definition of a term, it relates more

EXHIBIT F

THIS EXHIBIT HAS BEEN
REDACTED IN ITS ENTIRETY

EXHIBIT G

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT H

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT I

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY

EXHIBIT J

Main Identity

From: "Yochai Konig" <konig@speech.sri.com>
To: <stolcke@speech.sri.com>
Sent: Saturday, January 30, 1999 8:10 PM
Subject: 1998 Review

Hi Andreas,

I attach to this email my self-review for 1998. Please let me know if more information is need (e.g., papers) before our review meeting. Also please let me know when it would be convenient for you to meet.

Thanks,

--Yochai

1. 1998 Goals

My research in 1998 was centered around a data-driven approach for feature extraction for pattern recognition. By ``data-driven feature extraction" I mean the automatic extraction of features that optimize recognition performance. Feature extraction can be viewed as selecting a transformation from the original input space (e.g., digitized speech samples) to a smaller-dimension space. My approach of choosing this transformation according to recognition performance is in contrast to current feature extraction methods. This approach was applied to both speech recognition and speaker verification.

2. Speaker Verification

In speaker verification we study a nonlinear discriminant analysis (NLDA) technique that extracts a speaker-discriminant feature set. Our approach is to train a multilayer perceptron (MLP) to maximize the separation between speakers by nonlinearly projecting a large set of acoustic features (e.g., several frames) to a lower-dimensional feature set. The extracted features are optimized to discriminate between speakers and to be robust to mismatched training and testing conditions. We train the MLP on a development set and apply it to the training and testing utterances. Our results show that by combining the NLDA-based system with a state of the art cepstrum-based system we improve the speaker verification performance on the 1997 and 1998 NIST Speaker Recognition Evaluation sets by 15\% in average compared with our cepstrum-only system.

Achievements & Papers:

A. Part of the SRI team which officially won the 1998 speaker ID evaluations sponsored by NSA.

B. Publication and presentation at: Proc. RLA2C-ECSA, Speaker Recognition and its commercial and forensic applications, Avignon, France, April, 1998

- Konig, Y., Heck, L., Weintraub, M., and Sonmez, K., ``Nonlinear Discriminant Feature Extraction for Robust Text-Independent Speaker Recognition"

- Heck, L., and Konig, Y., ``Discriminative Training of Minimum Cost Speaker Verification Systems"

C. Submitted (with Larry, Kemal and Mitch) an extended version of `Nonlinear Discriminant Feature Extraction for Robust Text-Independent Speaker Recognition" paper to a special issue of Speech Communication magazine, following a selection process based on the best papers in the RLA2C conference. The two papers mentioned above were selected as a combined paper.

4. Speech Recognition

(Joint work with Mitch and Francoise)

- Approach

Our approach is to optimize all system components to maximize the posterior probability of the correct sentence. Our emphasis in this study is on the selection and estimation of the front-end model according to recognition performance. We optimize sentence-level measures and not frame-level measures. Specifically, we optimize the feature extraction process to increase the posterior probability of the correct sentence or of a specific cost function in case of a different error metric than word error rate (WER). We search for the optimal transformation from primitive features (e.g., FFT) to input features to HMM/GMM.

- Status

We derived and implemented an LVCSR system which jointly optimizes front-end and acoustic model according to sentence level criterion. This is the first time that a sentence level criterion was applied in a LVCSR system for optimizing front-end parameters. We explored research issues such as optimization criterion, MLP input features, GMM organization, Batch vs. Stochastic procedure. We obtained good performance for gender classification. However we achieved only modest gains for LVCSR with low capacity system - 1 Gaussian per class. Currently we continue this work in the context of the Marines

database. In addition, we extended this work to context dependent models.

- Publications

A. Beaufays, F., Weintraub, M., and Konig, Y., ``DYNAMO: An algorithm for Dynamic Acoustic Modeling",
In Proc. Broadcast News Transcription and Understanding Workshop (BNTU), Landsdowne, VA, February, 1998.

B. Beaufays, F., Weintraub, M., Konig, Y.,
``Discriminative Mixture Weight Estimation for Large Gaussian Mixture Models", to appear in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Tempa, AZ, March, 1999.

5. Projects - proposal writing

A. Proposed (as a PI) and awarded $200k for the IDEAS project based on the data - driven idea.

B. Took part in the LVCSR proposal (wrote two sections), joint work with Mitch, Andreas, and Francoise.

6. 1999 Goals

A. Contínting and extending my work on data-driven feature extraction both for speech recognition and speaker verification.

B. Exploring and studying statistical-based speech understanding.

EXHIBIT K

Main Identity

From: "Yochai Konig" <konig@speech.sri.com>
To: <mw@speech.sri.com>
Sent: Monday, June 08, 1998 10:56 AM
Subject: Ideas Proposal
*** EOOH ***
Return-Path: <konig@speech.sri.com>
Date: Mon, 8 Jun 1998 11:56:49 -0700
From: Yochai Konig <konig@speech.sri.com>
To: mw@speech.sri.com
Subject: Ideas Proposal

Hi Mitch,

Here is the draft for the ideas proposal.

--Yochai

DATA-DRIVEN FEATURE EXTRACTION
Yochai Konig and Mitch Weintraub
STAR LAB
==========

Introduction
==========

This draft proposal for the IDEAS program describes a data-driven approach for feature extraction for pattern recognition. The proposed approach has the potential to lead to major improvements in both speech recognition and speaker recognition performance as well as for other pattern recognition applications. In the speech recognition field feature extraction techniques such as vocal tract normalization (VTL), and optimized front-ends have all led to significant improvements as reported in the recent Large Vocabulary Continuous Speech Recognition (LVCSR) meetings by BBN, Dragon and SRI [LVCSR_nov97,LVCSR_may97]. In speaker recognition our work on data-driven features resulted in extracted features that are optimized to discriminate among speakers and to be robust to mismatched training and testing conditions. These features significantly improve speaker verification performance on the 1997 NIST Speaker Recognition Evaluation set compared with our cepstrum-only system \cite{konig_rla2c} and led to SRI officially winning the 1998 speaker ID evaluations sponsored by NSA.

We propose a data-driven approach for feature extraction for pattern

recognition. What we mean by ``data-driven feature extraction'' is the automatic extraction of features that optimize recognition performance. Feature extraction can be viewed as selecting a transformation from the original input space (e.g., digitized speech samples) to a smaller-dimension space. Our approach of choosing this transformation according to recognition performance is in sharp contrast to current feature extraction methods. Current techniques of feature extraction for speech recognition are knowledge-based methods that are based on auditory models (for example [seneff1986,Seneff88,ghitza1987]), perception theories [hermansky90], and signal processing considerations. The linkage to recognition performance is achieved by training systems with the proposed features and experimentally setting the technique's parameters.

The proposed feature extractor is data-driven in the sense that the extractor parameters are automatically estimated from (development) data, rather than being constrained to perform a specific function (e.g., spectral analysis). The advantages of the data-driven approach over traditional approaches which define an algorithm that compute a specific feature and than performing a trial and error of evaluating performance on a development set for each new set of features:

1. The search space for the optimal features is not constrained to a hand-tuned specific function which can lead to more discriminant features, thus, improved performance.

2. The process is automatic in contrast to the manual tedious trail and error process of optimizing current feature extraction methods.

The need for better and more discriminant features is driven by the overlapping distributions of the current models. This was, for example demonstrated by BBN in the May 1996 workshop [LVCSR_may96]. BBN reported in that meeting that the number of mixture components that contribute significant mass to a frame's probability is large and that even the simplest (triphone - state) distributions tend to cover a significant portion of the space given enough training material. However current techniques for optimizing front-ends and features are inefficient, and usually involve a trial and error process. Several researchers have suggested the incorporation of data-driven ideas in the feature extraction process. Rahim, Bengio, and LeCun suggested optimizing a set of parallel class-specific (e.g., phones) networks performing feature transformation based on the

minimum classification (MCE) criterion for telephone-based connected digit recognition [rahim97]. Fontaine, Ris, and Boite used the two-hidden-layer multilayer perceptron (MLP) to perform nonlinear discriminant analysis (NLDA) for isolated-word, large-vocabulary speech recognition tasks [fontaine97]. The training criterion for the MLPs was phonetic classification. Bengio and his colleagues suggested a global optimization of a neural network-hidden Markov (HMM) hybrid, where the outputs of the neural network constitute the observation sequence for the HMM [bengio92].

Approach

For speech recognition the novelty in our proposed approach is that we plan to go beyond cepstral representation for the input features to the feature extraction. The desired features will be extracted from local information such as the fast Fourier transform (FFT) points, global information such as speaking rate, and signal to noise ratio (SNR). We will take a large number of inputs and nonlinearly project them to a lower dimensional space based on a recognition criterion in common with some of the previous work. Our work with data-driven feature extraction methods works well for speaker verification tasks and shows potential for LVCSR tasks [konig_rla2c,LVCSR_nov97].

For speaker verification tasks, we have trained an MLP to maximize the separation between speakers by nonlinearly projecting a large set of acoustic features (e.g., several frames) to a lower-dimensional feature set. The extracted features are optimized to discriminate among speakers and to be robust to mismatched training and testing conditions. We train the MLP on a development set and apply it to the training and testing utterances. Our results show that by combining the system trained on the discriminantly extracted features, with a state-of-the-art cepstrum-based system, we improve speaker verification performance on the 1997 NIST Speaker Recognition Evaluation set by 15% in average compared with our cepstrum-only system [konig_rla2c]. However, our experiments with a similar approach for LVCSR tasks has yielded only marginal improvements so far. The main differences between the speaker-recognition tasks and the LVCSR experiments are that:

1. The training criterion for the feature extraction in the speaker verification task is the same as the overall performance measure, i.e., accuracy of speaker recognition. In the LVCSR experiments the training criterion for the feature extraction was single state phone discrimination, which differs from our overall goal of sentence recognition. Furthermore, the mismatch was even larger given that we used tristate triphones as our basic modeling unit and not single state phone as was the measure for the feature extraction.

2. In the speaker verification study we used other features in addition to

cepstrum as inputs to the feature extraction process (e.g., estimation of pitch). This did not occur in the LVCSR task.

Based on these differences we propose a new approach for data-driven feature extraction for LVCSR tasks. This new approach provides solutions as follows

1. To overcome the mismatch between the feature extraction training criterion and the overall recognition performance criterion in, we propose to optimize feature extraction according to recognition performance. We will optimize sentence level measures and not frame level measures. Specifically, we will optimize the feature extraction process to increase the posterior probability of the correct sentence or of a specific cost function in case of a different error metric than word error rate (WER).

2. We plan to go beyond the cepstrum for the input features to feature extraction. Specifically, we plan to use the fast FFT points themselves (same information as the speech samples). In addition we will use features that reflect global correlation of the test data speaker, dialect, and channel, and we will perform nonlinear dimension reduction based on recognition performance.

Research Issues

An important research question is what should be the inputs to the feature extraction process. We plan to use the FFT points (both the real and imaginary parts) for a large window of speech as inputs. By using FFT points we make no assumptions about the nature of the extracted features because FFT points carry exactly the same information as speech samples. We will initialize the MLP by training it to map from the FFT points to the cepstral features. The reasoning for using FFT points instead of speech samples as inputs to the MLP is that FFT points have internal repeatable order (as opposed to the waveform where a shift in time of several samples will drastically change the representation). We can train an MLP to approximate any function given enough training patterns and enough hidden units [neural_comp].

The input representation to the MLP can be augmented to (a) include other types of information, and (b) to make the representation more efficient. To include other types of information, we plan to augment the input to the MLP with longer term or global information. Information such as such as time derivatives of the cepstral parameters, speaking rate, VTL, hidden state variables (see section on long-term correlation modeling), and signal to noise ratio (SNR) can also be used as input features to the MLP.

To make the representation more efficient, we plan to study techniques to find minimal configurations of the feature extractor. The training

of the MLP to map from an input feature such as the FFT to a cepstral feature vector is feasible (since we can use an infinite amount of data to train this mapping). However, the unknown variable is the number of MLP parameters that are needed to perform this task. If the number of parameters is too large, this will make it difficult to move away from the initial MLP-implementation of the cepstral transformation with a limited amount of labeled data (the training speech corpus). Therefore, a critical part of the research is how to efficiently encode the discriminative information with a minimal number of parameters. One way to make the representation more efficient is to include additional knowledge sources as inputs to the MLP (e.g. FFT energies, original cepstral representation) as well as use algorithms for model selection. Model selection is well studied [brain_damage,moody_94]. However in the speech community the problem of automatic model selection based on recognition performance has not been extensively studied.

Another important research area is the interaction with model parameter estimation (e.g. HMM output distributions). We propose to study a joint optimization of the model parameters and the features at the frame level is better than an iterative procedure similar in nature to the expectation maximization (EM) [dempster77] Based on these research issues and to be concrete we outline two sample studies for our approach.

Sample Study 1

We plan to start from the FFT points for a large window of speech as inputs to our feature extractor. Initially, we will train an MLP to map these FFT points to a standard feature vector (e.g., 10 cepstral coefficients and their first and second time derivatives) for a good initialization point. In the second stage we will back-propagate the error into this MLP with the criterion of maximizing the posterior probability of the correct sentence, by using a stochastic gradient approach. After the training, we will then have a transformation of the FFT points into a new feature vector (i.e., the output of the MLP) that we can use to process our data. A natural augmentation for the input features is the incorporation of global input features that involve longer time correlations reflecting speaking rate, accent, and channel estimation.

Sample Study 2

We will adapt the model parameters (such as means and variances) in addition to the features, according to the same criterion of increasing the posterior probability of the correct sentence. We will experiment to determine whether the joint optimization of the model parameters and the features will be at the frame level or iterative in nature (similar to expectation maximization EM).

EXHIBIT L

THIS EXHIBIT HAS BEEN REDACTED IN ITS ENTIRETY