

Robust heteroscedastic linear discriminant analysis and LCRC posterior features in meeting data recognition

Martin Karafiát, František Grézl, Petr Schwarz, Lukáš Burget, and
Jan Černocký

Speech@FIT, Faculty of Information Technology, Brno University of Technology
{karafiat, grezl, schwarzp, burget, cernocky}@fit.vutbr.cz

Abstract. This paper investigates into feature extraction for meeting recognition. Three robust variants of popular HLDA transform are investigated. Influence of adding posterior features to PLP feature stream is studied. The experimental results are obtained on two data-sets: CTS (continuous telephone speech) and meeting data from NIST RT'05 evaluations. Silence-reduced HLDA and LCRC phoneme-state posterior features are found to be suitable for both recognition tasks.

1 Introduction

The AMI project¹ concentrates at processing and analysis of meetings. One of key problems is to determine what was said in the meeting; this task is accomplished by large vocabulary continuous speech recognition (LVCSR). This paper deals with the extraction of features from speech signal.

One of key problems in feature extraction is to reduce the dimensionality of feature vectors while preserving the discriminative power of features. Linear transforms such as Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA) are mostly used for this task. In recent years, Heteroscedastic Linear Discriminant Analysis (HLDA) has gained popularity in the research community [2, 1] for its relaxed constraints on statistical properties of classes (Unlike LDA, HLDA does not assume same covariance matrix for all classes). To compute HLDA transformation matrix, however, more statistics need to be estimated and the reliability of such estimations becomes an issue. Section 2 discusses robust variants of HLDA.

Second part of the paper is devoted to the use of posterior-features. Posteriors generated by neural networks (NN) and converted into features are also increasingly popular in small [7] and large [6] recognition systems for their complementarity with classical PLP or MFCC coefficients. Section 3 introduces phoneme-state posterior estimator based on split temporal context [8, 9] that has already proved its quality in different tasks ranging from language identification to keyword spotting.

¹ Augmented Multi-Party Interaction <http://www.amiproject.org>

Most of the development work was done on a continuous telephone speech system (section 4). For the recognition of meetings, we used NIST RT'05 data and took advantage of the AMI-LVCSR system [5]². The results described in section 5 are obtained by re-scoring LVCSR lattices generated by the AMI system.

2 HLDA

HLDA allows to derive such projection that best de-correlates features associated with each particular class (maximum likelihood linear transformation for diagonal covariance modeling [2]). To perform de-correlation and dimensionality reduction, n -dimensional feature vectors are projected into first $p < n$ rows, $\mathbf{a}_{k=1\dots p}$, of $n \times n$ HLDA transformation matrix, \mathbf{A} . An efficient iterative algorithm [3, 1] is used in our experiments to estimate matrix \mathbf{A} , where individual rows are periodically re-estimated using the following formula:

$$\hat{\mathbf{a}}_k = \mathbf{c}_k \mathbf{G}^{(k)-1} \sqrt{\frac{T}{\mathbf{c}_k \mathbf{G}^{(k)-1} \mathbf{c}_k^T}} \quad (1)$$

where \mathbf{c}_i is the i^{th} row vector of co-factor matrix $C = |\mathbf{A}| \mathbf{A}^{-1}$ for current estimate of \mathbf{A} and

$$\mathbf{G}^{(k)} = \begin{cases} \sum_{j=1}^J \frac{\gamma_j}{\mathbf{a}_k \hat{\Sigma}^{(j)} \mathbf{a}_k^T} \hat{\Sigma}^{(j)} & k \leq p \\ \frac{T}{\mathbf{a}_k \hat{\Sigma} \mathbf{a}_k^T} \hat{\Sigma} & k > p \end{cases} \quad (2)$$

where $\hat{\Sigma}$ and $\hat{\Sigma}^{(j)}$ are estimates of global covariance matrix and covariance matrix of j^{th} class, γ_j is number of training feature vectors belonging to j^{th} class and T is the total number of training feature vectors. In our experiments, the classes are defined by each Gaussian mixture component and γ_j are their occupation counts.

Well known Linear Discriminant Analysis (LDA) can be seen as special case of HLDA, where it is assumed that covariance matrices of all classes are the same. In contrast to HLDA, closed form solution exists in this case. Basis of LDA transformation are given by eigen vectors of matrix $\Sigma_{AC} \times \Sigma_{WC}^{-1}$, where Σ_{WC} is within-class covariance matrix and Σ_{AC} is across-class covariance matrix.

2.1 SHLDA

HLDA requires the covariance matrix to be estimated for each class. The higher number of classes is used, the fewer feature vector examples are available for each class and class covariance matrix estimates become more noisy. We have

² Brno University of Technology is a member of AMI-LVCSR development team, coordinated by University of Sheffield.

recently proposed [1] a technique based on combination of HLDA and LDA, where class covariance matrices are estimated more robustly, and at the same time, (at least the major) differences between covariance matrices of different classes are preserved. Smoothed HLDA (SHLDA) differs from HLDA only in the way of class covariance matrices estimation. In the case of SHLDA, estimate of class covariance matrices is given by:

$$\tilde{\Sigma}_j = \alpha \hat{\Sigma}_j + (1 - \alpha) \Sigma_{WC} \quad (3)$$

where $\tilde{\Sigma}_j$ is “smoothed” estimate of covariance matrix for class j . $\hat{\Sigma}_j$ is estimate of covariance matrix, Σ_{WC} is estimate of within-class covariance matrix and α is smoothing factor — a value in the range of 0 to 1. Note that for α equal to 0, SHLDA becomes LDA and for α equal to 1, SHLDA becomes HLDA.

2.2 MAP-SHLDA

SHLDA gives more robust estimation than standard HLDA but optimal smoothing factor α depends on the amount of data for each class. In extreme case, α should be set to 0 (HLDA) if infinite amount of training data is available. With decreasing amount of data, optimal α value will slide up to LDA direction.

To add more robustness into the smoothing procedure, we implemented maximum a posteriori (MAP) smoothing [4], where within-class covariance matrix Σ_{WC} is considered as the prior. Estimate of the class covariance matrix is then given by:

$$\tilde{\Sigma}_j = \Sigma_{WC} \frac{\tau}{\gamma_j + \tau} + \hat{\Sigma}_j \frac{\gamma_j}{\gamma_j + \tau} \quad (4)$$

where τ is a control constant. Obviously, if insufficient data is available for current class, the prior resource Σ_{WC} is considered as more reliable than the class estimation $\hat{\Sigma}_j$. In case of infinite data, only the class estimation of covariance matrix $\hat{\Sigma}_j$ is used for further processing.

2.3 Silence Reduction in HLDA

From the point of view of transformation estimation, silence is a “bad” class as its distributions differ significantly from all speech classes. Moreover, training data (even if end-pointed) contains significant proportion of silence. Therefore, we have experimented with limiting the influence of silence.

Rather than discarding the silence frames, the occupation counts, γ_j , of silence classes, which takes part in computation of global covariance matrix, $\hat{\Sigma}$, and in equation 2 are scaled by factor $1/SR$

$SR = \infty$ corresponds to complete elimination of silence statistics.

3 Posterior features

Several works have shown that using posterior-features generated by NNs is advantageous for speech recognition [7, 6]. We have experimented with two setups




Fig. 1. Phoneme-state posterior estimator based on split left and right contexts.

to generate posteriors. The first one is based on a simple estimation of phoneme posterior probabilities from a block of 9 consecutive PLP-feature vectors (FeatureNet).

The second one uses our state-of-the-art phoneme-state posterior estimator based on modeling long temporal context[9]. Details of the posterior estimator are shown in Fig. 1. Mel filter bank log energies are obtained in conventional way. Based on our previous work in phoneme recognition [8], the context of 31 frames (310 ms) around the current frame is taken. This context is split into 2 halves: Left and Right Contexts (hence the name “LCRC”). This allows for more precise modeling of the whole trajectory while limiting the size of the model (number of weights in the NN) and reducing the amount of necessary training data. For both parts, temporal evolutions of critical band log energies are processed by discrete cosine transform to de-correlate and reduce dimensionality. Two NNs are trained to produce phoneme-state posterior probabilities for both context parts. We use 3 states per phoneme which follows similar idea as states in phoneme HMM. Third NN functions as a merger and produces final set of phoneme-state posterior probabilities³

For both approaches, the resulting posteriors are processed by log and by a linear transform to de-correlate and reduce dimensionality (details are given in experimental sections 4 and 5).

4 CTS experiments

Our recognition system was trained on ctstrain04 training set, a subset of the h5train03 set, defined at the Cambridge University as a training set for Conversation Telephone Speech (CTS) recognition systems [5]. It contains about 278 hours of well transcribed speech data from Switchboard I, II and Call Home English. All systems were tested on the Hub5 Eval01 test set composed of 3 subsets of 20 conversations from Switchboard I, II, and Switchboard-cellular corpora, for a total length of about 6 hours of audio data.

The baseline features are 13th order PLP cepstral coefficients, including 0th one, with first and second derivatives added. This gives a standard 39 dimension

³ Neural nets are trained using QuickNet from ICSI and SNet – a parallel NN training software being developed in Speech@FIT

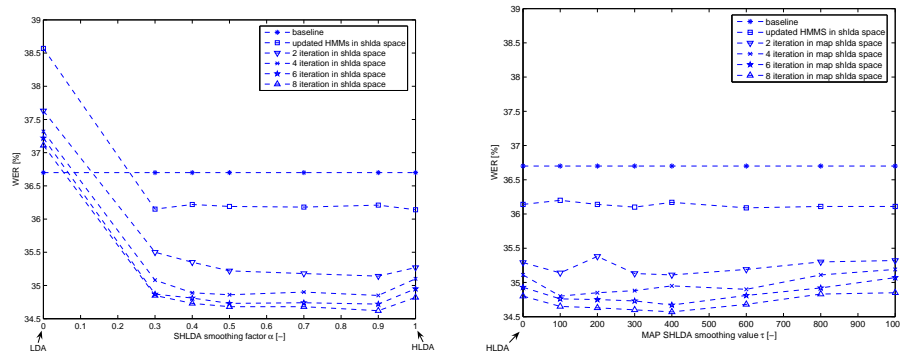


Fig. 2. Dependency of WER on the SHLDA (left) and MAP-SHLDA (right) smoothing factors.

feature vector. Cepstral mean and variance normalization was applied. Baseline cross-word triphone HMM models were trained by Baum-Welch re-estimation and mixture splitting. We used a standard 3-state left-to-right phoneme setup, with 16 Gaussian mixtures per state. 7598 tied states were obtained by decision tree clustering. Each Gaussian mixture was taken as a different class for HLDA experiment. Therefore, we had $N = 16 \times 7598 = 121568$ classes.

The tri-gram language model used in decoding setup was computed by interpolation from Switchboard I, II, Call Home English and Hub4 (Broadcast news) transcriptions. The size of recognition vocabulary was 50k words.

The recognition output was generated in two passes: At first, lattice generation with baseline HMMs and bigram language model was performed. The lattices were expanded by more accurate trigram language model. The pruning process was applied to reduce them to reasonable size. In the second pass, lattices were re-scored with tested features and models.

4.1 Flavors of HLDA

We added the third derivatives into the feature stream, which gave us 52 dimensional feature vectors. **SHLDA** transform was then trained to perform the projection from 52 to 39 dimension. Smoothing factors α in Eq. 3 of 0.0 (LDA), 0.3, 0.4, 0.5, 0.7, 0.9, 1.0 (HLDA) were tested. Left panel of figure 2 shows dependency of WER on SHLDA smoothing factor α . Pure LDA failed, probably due to bad assumption of the same Gaussian distribution in all classes. The best system performance (Table 1) was obtained for smoothing factor 0.9. The relative improvement of this system is 7.9% compared to the baseline and 0.6% compared to the clean HLDA setup.

MAP-SHLDA test setup was built in same way as SHLDA system, only the smoothing procedure (Equation 3) was replaced by MAP approach (Equation 4). The average value of all class occupation counts was 820. Therefore $\tau = 820$

System	WER [%]
Baseline (no HLDA)	36.7
HLDA	34.8
SHLDA	34.6
MAP-SHLDA	34.6
SR-HLDA	34.5

Table 1. Comparison of HLDA systems.

in MAP-SHLDA should have the same behavior as $\alpha = 0.5$ in SHLDA if all classes had the same number of observations. The optimal smoothing values for SHLDA were in range 0.5—0.9 (left panel of Figure 2). Therefore, we decided to test optimal smoothing control constant τ on values 0 (HLDA), 100, 200, 300, 400, 600, 800 and 1000. The results are shown in right panel of Figure 2. The best system performance (Table 1) was obtained for $\tau = 400$. The relative improvement of this system is 8% compared to the baseline and 0.7% compared to the clean HLDA setup.

Silence reduction in HLDA (SR-HLDA) was tested with factors SR equal to 1 (no reduction), 2, 10, 100 and ∞ (removing all silence classes). For $SR = 1$, the WER is obviously 34.8%, for $SR = 2$ it drops to 34.6% and from $SR = 10 \dots \infty$ it is constant: 34.5%.

4.2 Posterior features

Posterior features were always used together with base PLP features. Table 2 summarizes the results.

Upper part of Figure 3 shows the way the two feature streams were combined in FeatureNet experiments. The upper branch corresponds to the previous section. To compute posterior features, 9 frames of PLP+ Δ + $\Delta\Delta$ were stacked and processed by a neural net with 1262 neurons in the hidden layer (this number was chosen to have approximately 500k weights in the NN). There are 45 phoneme classes, which determines the size of the output layer. Log-posteriors are processed by KLT or HLDA and then concatenated with PLP+HLDA features to form the final 64-dimensional feature vectors.

Lower panel of Figure 3 presents the setup with LCRC-posterior features. The PLPs were derived directly with Δ , $\Delta\Delta$ and $\Delta\Delta\Delta$, and down-scaled by HLDA to 39 dimensions. The detail of LCRC-posterior feature derivation is in Fig. 1, all nets had 1500 neurons in the hidden layer. For each frame, the output of LCRC system are estimates of 135 phoneme-state⁴ posterior probabilities. As the number of phoneme-state posteriors is too high to fit the statistics necessary for HLDA estimation into the memory, the output dimensionality of LCRC system is first reduced by KLT from 135 to 70. The following HLDA reduces this size to

⁴ see [9] for details on splitting each of phonemes to 3 phoneme-states

System	WER [%]
PLP SR-HLDA	34.5
PLP SR-HLDA + PLP-posteriors KLT	33.8
PLP SR-HLDA + PLP-posteriors HLDA	33.3
PLP SR-HLDA + LCRC-posteriors HLDA	32.6

Table 2. Performance of posterior features in the CTS system.

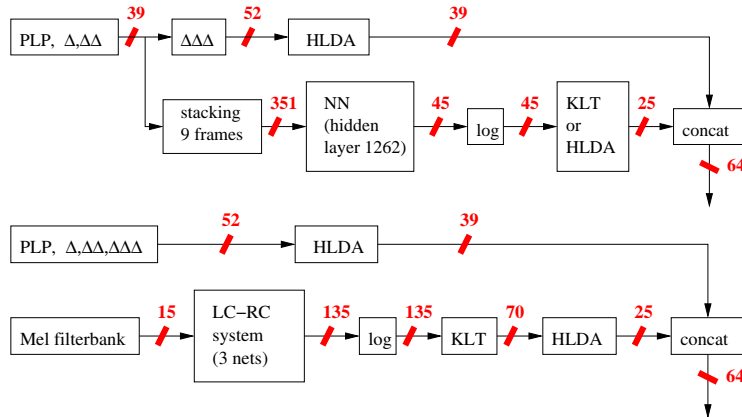


Fig. 3. Configuration of the system with PLP- (upper panel) LCRC-posteriors (lower panel).

25, and the results are concatenated with PLP+HLDA features to form again 64-dimensional feature vectors.

We see, that the posterior features improve the results by almost 1% absolutely, and that there is clear preference of HLDA to KLT. With the new LCRC features, we have confirmed good results they provide in phoneme recognition [9] — with these features, the results are almost 2% better than the PLP SR-HLDA baseline.

5 Meeting data experiments

Training and test sets for these experiments are the same as those used by AMI-LVCSR system for NIST RT'05 evaluations [5], therefore, we limit ourselves to the most important details. The training data consists of more than 100 hours of meeting data originating from ICSI meeting corpus (73h), NIST data (13h), ISL (10h) and AMI preliminary development set (16h).

The test data comes from NIST RT'05 and consists of two 10-minute excerpts from meetings collected by ICSI, NIST, ISL, AMI, and Virginia Polytechnic and State University (VT). NIST RT'05 included audio from headset microphones

System	WER [%]
PLP SR-HLDA	28.7
PLP SR-HLDA + LCRC-posteriors HLDA	26.0

Table 3. Performance of HLDA and posterior features in meeting recognition.

(Independent Headset Microphone, IHM) and from table-top microphones (Multiple Distant Microphones, MDM), in this work, only IHM condition was used.

The system architecture is described in detail in [5]. The system operates in a total of 6 passes according to Fig. 4. The 4-gram lattices generated in step P4 were taken as input for the tests described in this section and were re-scored with different feature-extraction setups.

5.1 HLDA and posterior features in meeting recognition

Table 3 shows the results for different feature setups. The baseline for these experiments are PLP features with VTLN (VTLN was applied also prior to any posterior feature derivation) and SR-HLDA for de-correlation and dimensionality reduction (PLP $\Delta\Delta\Delta$ (52 dimensions) \rightarrow 39). SR-HLDA was selected for good and “cheap” performance in the CTS system.

Posterior features were generated with the LCRC-system accordingly to Fig. 1 and lower panel of Fig. 3. The only changes from the CTS setup were the use of 23 filters in the bank instead of 15 (wide-band speech) and consistent application of VTLN prior to both PLP and Mel-filterbank computation. The sizes of hidden layers in neural nets were the same (1500 neurons), the size of input layers in left- and right-context nets increased due to increased numbers of bands. The improvement obtained by LCRC-posterior features is again quite impressive – 2.7%.

6 Conclusion

In this paper, we have investigated robust variants of HLDA and use of classical and novel posterior features in telephone speech and meeting data recognition.

In the HLDA part, 2 approaches of HLDA smoothing were tested: Smoothed HLDA (SHLDA) and MAP variant of SHLDA taking into account the amounts of data available for estimation of statistics for different classes. Both perform better than the basic HLDA. We have however found, that removing of silence class from the HLDA estimations (Silence-reduced HLDA) is equally effective and cheaper in computation. Testing SHLDA and MAP-SHLDA on the top of SR-HLDA did not bring any further improvement, therefore we stick with SR-HLDA as the most suitable transformation in our LVCSR experiments.

Two kinds of posterior features were tested – “classical” FeatureNet approach with stacked 9 frames of PLPs and novel approach using more elaborate structure to phoneme-state posterior modeling. The later scheme provided significant reduction of error rate in both CTS and meeting data experiments.

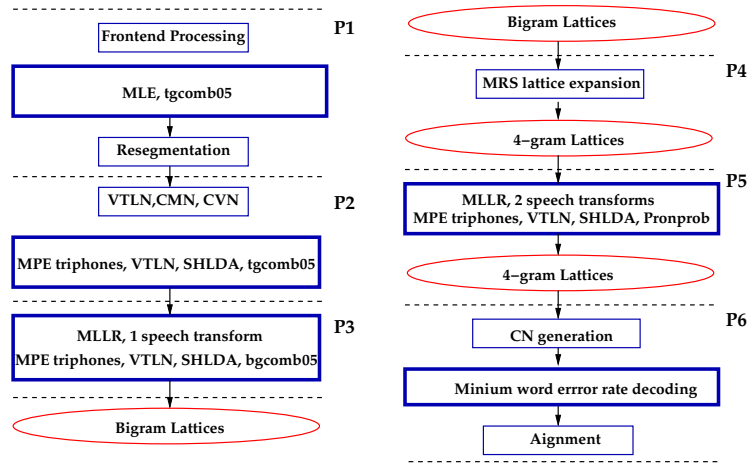


Fig. 4. Processing stages of AMI system for 2005 NIST evaluations. The figure is reprinted from [5] with permission of the author.

In our future work, we will investigate if the improvement obtained by LCRC-posteriors is preserved after discriminative training and speaker adaptation (MLLR) applied on the top of such features. Preliminary results are quite promising. The described features will be integrated into AMI-LVCSR system in NIST RT'06 evaluations.

7 Acknowledgments

This work was partially supported by EC project Augmented Multi-party Interaction (AMI), No. 506811 and Grant Agency of Czech Republic under project No. 102/05/0278. Lukáš Burget was supported by post-doctoral grant of Grant Agency of Czech Republic No. 102/06/P383. Thanks University of Sheffield for generating LVCSR lattices. We further thank Cambridge University Engineering Department making the h5train03 CTS training set available for granting the right to use Gunnar Evermann's HDecode to the University of Sheffield.

References

1. L. Burget, "Combination of speech features using smoothed heteroscedastic linear discriminant analysis," in *8th International Conference on Spoken Language Processing*, Jeju island, KR, oct 2004.
2. N. Kumar, *Investigation of Silicon-Auditory Models and Generalization of Linear Discriminant Analysis for Improved Speech Recognition*, Ph.D. thesis, John Hopkins University, Baltimore, 1997.
3. M.J.F. Gales., "Semi-tied covariance matrices for hidden markov models," *IEEE Trans. Speech and Audio Processing*, vol. 7, pp. 272–281, 1999.
4. J. Gauvain and C. Lee, "Maximum a posteriori estimation for multivariate Gaussian mixture observations of Markov chains," *IEEE Trans. Speech Audio Processing*, vol. 2, pp. 291–298, Apr. 1994.
5. T. Hain et al., "The 2005 AMI system for the transcription of speech in meetings," in *Proc. Rich Transcription 2005 Spring Meeting Recognition Evaluation Workshop*, Edinburgh, July 2005.
6. Q. Zhu, A. Stolcke, B. Y. Chen and N. Morgan: "Using MLP Features in SRI's Conversational Speech Recognition System" in *Proc. Eurospeech 2005*, pp. 2141–2144, Lisabon, Portugal, 2005.
7. A. Adami, L. Burget, S. Dupont, H. Garudadri, F. Grezl, H. Hermansky, P. Jain, S. Kajarekar, N. Morgan, and S. Sivasdas. Qualcomm-ICSI-OGI features for ASR. In *Proc. ICSLP 2002*, Denver, Colorado, USA, 2002.
8. P. Schwarz, P. Matějka, and J. Černocký, "Towards lower error rates in phoneme recognition," in *Proc. International Conference on Text, Speech and Dialogue*, Brno, Czech Republic, Sept. 2004, pp. 465–472.
9. Petr Schwarz, Pavel Matějka and Jan Černocký: "Hierarchical structures of neural networks for phoneme recognition", accepted to ICASSP 2006, Toulouse, 2006.