

FEATURE ADAPTATION BASED ON GAUSSIAN POSTERIORIS

Suleyman S. Kozat, Karthik Visweswariah and Ramesh Gopinath

IBM T.J. Watson Research Center
Yorktown Heights, NY 10598

{kozat, kv1, rameshg}@us.ibm.com

ABSTRACT

In this paper we consider the use of non-linear methods for feature adaptation to reduce the mismatch between test and training conditions. The non-linearity is introduced by using the posteriors of a set of Gaussians to (softly) partition the observation space for feature adaptation. The modeling framework used is based on the fMPE models [1] applied to FMLLR matrices directly. However, the parameters are estimated to maximize the likelihood of the test data. We observe a relative gain of 14% on top of FMLLR, which was a 42% relative gain over the baseline.

1. INTRODUCTION

State of the art speech recognitions systems typically adapt their features and/or acoustic models to the test speaker to get improved recognition accuracy. In this paper we only consider adapting the features. Popular techniques for feature adaptation/normalization include spectral subtraction, Code-word Dependent Cepstral Normalization (CDCN) [2] and Feature space Maximum Likelihood Linear Regression (FMLLR) [3]. FMLLR is a linear technique where the features are linearly transformed to maximize the likelihood of the test data under a given fixed model. FMLLR differs from most of the other techniques in that no explicit assumptions are made about the type of noise or channel. Although FMLLR has been quite successful, several attempts have been made at generalizing the technique to allow for non-linear transforms of the feature vectors [4], [5]. [4] and [6] consider non-linear transforms at training time.

In this paper, we present a nonlinear extension of FMLLR by using Gaussian posteriors which has roots in fMPE [1]. Unlike standard FMLLR [3] which fits a single transformation matrix to the observation model, we partition the observation model into regions and estimate a separate FMLLR transformation for each region, as seen in Figure 1. Each feature vector is assigned to these regions using aposterior probabilities calculated by a Gaussian Mixture Model (GMM). The final transformation for a particular feature vector is constructed as the combination of all the transformation matrices weighted by the appropriate aposterior probability. This




Fig. 1. a) USING a single FMLLR transformation for the observation model b) Using multiple FMLLR matrices for the observation model, for each region separately.

method can be considered as the piecewise linear extension of the standard FMLLR method. The decision boundaries of each region is “soft” since we use the aposterior occupation probabilities in calculation of the final transform. Although we use the posteriors from a GMM to introduce the non-linearity, the methods used to estimate parameters will be independent of the actual non-linearity used.

The rest of this paper is organized as follows. In Section 2 we describe the feature transformation model and the objective we use to estimate the parameters. In Section 3 we present the technique used to estimate parameters. Section 4.2 describes the databases and the experimental setup used to evaluate our techniques, and presents our results on this database. We present our conclusions and some directions for future work in Section 5.

2. FEATURE TRANSFORMATION MODEL AND OBJECTIVE FUNCTION

Let us denote the feature vector at time t by \vec{x}_t . Then the basic model we use to generate the transformed feature \vec{y}_t is

$$\begin{aligned}\vec{y}_t &= (\mathbf{A}_1\phi_1(\vec{x}_t) + \mathbf{A}_2\phi_2(\vec{x}_t) + \dots + \mathbf{A}_m\phi_m(\vec{x}_t))\vec{x}_t, \quad (1) \\ &= \sum_{g=1}^m \mathbf{A}_g\phi_g(\vec{x}_t)\vec{x}_t,\end{aligned}$$

where each \mathbf{A}_g is an FMLLR matrix of appropriate dimension. Although the estimation techniques apply to general

$\vec{\phi}(\vec{x}_t) \stackrel{\text{def}}{=} [\phi_1(\vec{x}_t), \dots, \phi_m(\vec{x}_t)]^T$, in this paper we only consider the use of Gaussian posteriors [1]. We assume we have a given fixed GMM with m Gaussians which we use to calculate the g th component of $\phi(\vec{x}_t)$ as:

$$\phi_g(\vec{x}_t) = \frac{\pi_g N(\vec{x}_t; \vec{\mu}_g, \Sigma_g)}{\sum_k \pi_k N(\vec{x}_t; \vec{\mu}_k, \Sigma_k)},$$

where $N(\vec{x}; \vec{\mu}, \Sigma)$ denotes the likelihood of \vec{x} under a Gaussian density with mean $\vec{\mu}$ and covariance Σ .

For the optimization purposes, we can write Equation (1) as

$$\vec{y}_t = \mathbf{A} \vec{f}(\vec{x}_t)$$

where $\mathbf{A} \stackrel{\text{def}}{=} [\mathbf{A}_1 \dots \mathbf{A}_m]$ and

$$\vec{f}(\vec{x}_t) \stackrel{\text{def}}{=} \begin{bmatrix} \phi_1(\vec{x}_t) \vec{x}_t \\ \vdots \\ \phi_m(\vec{x}_t) \vec{x}_t \end{bmatrix}.$$

We would like to estimate our parameters \mathbf{A} to maximize the likelihood of the test data. For this to be valid we need to ensure that the feature transform we are using is invertible and compensate the likelihood with the log determinant of the Jacobian. Let M denote the GMM and G denote a graph which specifies a set of allowed state sequence. Then the objective function we need to maximize is:

$$g(\mathbf{A}) = \log \left| \det \frac{d\mathbf{Y}}{d\mathbf{X}} \right| + \log P(\mathbf{Y}|M, G),$$

where \mathbf{X} denotes all the acoustic features for a particular test speaker and \mathbf{Y} denotes the transformed acoustic features for that speaker. We split this into the Jacobian term and the likelihood term which are handled differently:

$$g_J(\mathbf{A}) = \log \left| \det \frac{d\mathbf{Y}}{d\mathbf{X}} \right|$$

and $g_L(\mathbf{A}) = \log P(\mathbf{Y}|M, G)$. Note that since \vec{y}_t is a function of only \vec{x}_t ,

$$\log \left| \det \frac{d\mathbf{Y}}{d\mathbf{X}} \right| = \sum_t \log \left| \det \frac{d\vec{y}_t}{d\vec{x}_t} \right|.$$

Ensuring that our transform is invertible is equivalent to ensuring that the Jacobian is full rank for all \mathbf{X} . This is hard to do in general, and we do not deal with the issue rigorously. We note that the log determinant term in the objective function goes to negative infinity when the Jacobian becomes singular. We assume this will prevent us, in practice, from making the transform non-invertible.

3. PARAMETER ESTIMATION

We use the limited memory BFGS algorithm [7] with the More-Thuente line search algorithm [8] as implemented in [9]

to minimize g . This requires computation of g and its gradient with respect to \mathbf{A} . Each computation of g requires a pass through the adaptation data. Note that we do not use an auxiliary function to optimize the objective function. Using an auxiliary function does not give us the usual benefit of being able to go through the data once and collect sufficient statistics, which can be used to perform the optimization. This is because of the Jacobian term g_J , for which we need to run through the data each time we want to calculate it and its gradient.

Let us now go into the calculation of g_L and its gradient. First we note that if we can calculate gradient g_L w.r.t \vec{y} then we can propagate this gradient using the chain rule to calculate all required gradients as follows:

$$\frac{dg_L}{d\mathbf{A}} = \frac{dg_L}{d\mathbf{Y}} \frac{d\mathbf{Y}}{d\mathbf{A}}. \quad (2)$$

Let \mathcal{G} be the set of Gaussian sequences determined by the model M and the graph G . Then we can write:

$$g_L = \log P(\mathbf{Y}|M, G) = \log \sum_{g^n \in \mathcal{G}} P(g^n) P(\mathbf{Y}|g^n).$$

The gradient $g_L(\vec{y})$ w.r.t a given frame \vec{y}_t is given by:

$$\begin{aligned} \frac{d \log P(\mathbf{Y}|M, G)}{d\vec{y}_t} &= \frac{d \log \sum_{g^n \in \mathcal{G}} P(g^n) P(\mathbf{Y}|g^n)}{d\vec{y}_t} \\ &= \frac{\sum_{g^n \in \mathcal{G}} P(g^n) dP(\mathbf{Y}|g^n)/d\vec{y}_t}{\sum_{g^n \in \mathcal{G}} P(g^n) P(\mathbf{Y}|g^n)} \\ &= \sum_{g \in \mathcal{G}_t} \gamma_g(t) \frac{d \log P(\vec{y}_t|g)}{d\vec{y}_t} \\ &= \sum_{g \in \mathcal{G}_t} \gamma_g(t) \Sigma_g^{-1} (\vec{\mu}_g - \vec{y}_t), \end{aligned}$$

where \mathcal{G}_t is the set of Gaussians that are allowed at time t according to the set of Gaussian sequences \mathcal{G} . Plugging this result into Equation 2, we get

$$\frac{dg_L}{d\mathbf{A}} = \sum_t \sum_{g \in \mathcal{G}_t} \gamma_g(t) \Sigma_g^{-1} (\vec{\mu}_g - \vec{y}_t) \vec{f}(\vec{x}_t)^T. \quad (3)$$

We now turn to the calculation of g_J and it's gradient. The Jacobian of our feature transform is

$$g_J = \sum_t \log \left| \det \frac{d\vec{y}_t}{d\vec{x}_t} \right| = \sum_t \log \left| \det \left(\mathbf{A} \frac{d\vec{f}(\vec{x}_t)}{d\vec{x}_t} \right) \right|.$$

It can be shown that,

$$\frac{d\vec{f}(\vec{x}_t)}{d\vec{x}_t} = \begin{bmatrix} \vec{x}_t \left(\frac{d\phi_1(\vec{x}_t)}{d\vec{x}_t} \right)^T + \phi_1(\vec{x}_t) \mathbf{I} \\ \vdots \\ \vec{x}_t \left(\frac{d\phi_m(\vec{x}_t)}{d\vec{x}_t} \right)^T + \phi_m(\vec{x}_t) \mathbf{I} \end{bmatrix}$$

where \mathbf{I} is an appropriate size identity matrix.

Let us consider $\phi_g(\vec{x}_t)$ the g th component of $\phi(\vec{x}_t)$. The derivative of this is:

$$\frac{d\phi_g(\vec{x}_t)}{d\vec{x}_t} = \phi_g(\vec{x}_t) \left(\Sigma_g^{-1}(\vec{\mu}_g - \vec{x}_t) - \sum_k \phi_k(\vec{x}_t) \Sigma_k^{-1}(\vec{\mu}_k - \vec{x}_t) \right).$$

Note that this gradient is zero when $\phi_g(\vec{x}_t) = 0$ or $\phi_g(\vec{x}_t) = 1$. Thus the gradient of g_J is

$$\frac{dg_J}{d\mathbf{A}} = \sum_t \left(\mathbf{A} \frac{d\vec{f}(\vec{x}_t)}{d\vec{x}_t} \right)^{-T} \frac{d\vec{f}(\vec{x}_t)}{d\vec{x}_t}^T.$$

4. EXPERIMENTAL SETUP AND RESULTS

4.1. Training and test database description

The experiments reported on in this paper were performed on an IBM internal database [10]. The test data consists of utterances recorded in a car at three different speeds: idling, 30 mph and 60 mph. Four tasks are included in the test set: addresses, digits, commands and radio control. Following are typical utterances from each task:

A: *New York City ninety sixth street West.*

C: *Set track number to seven.*

D: *Nine three two three three zero zero.*

The test data base has 73743 words and each speaker has on the average 5.2 minutes of data.

Training data was also collected in a car at three speeds. Since most of the data was collected in a stationary car, the training data was augmented by adding noise collected in a car to the data collected in a stationary car. Data was collected with microphones in three different positions: rear-view mirror, visor and seat belt. The database used for training consisted of 887110 utterances. The baseline acoustic model was word internal with 826 states and 10001 diagonal Gaussians. The front end we use is fairly standard; MFCC (13 dimensional) with mean normalization (max normalization for c0) and delta and double deltas (final feature 39 dimensional).

All of our experiments are unsupervised adaptation experiments. We first decode the test data, and then used the decoded script to generate a forced alignment. This alignment is then used as the graph G used to calculate the likelihood g_L . We could use the decoding graph or the decode word script instead but past experience shows that this is usually no better than using an alignment of the decoded script.

4.2. Results

The baseline error rate on our test database is 2.08%. The standard FMLLR gives an error rate of 1.41%. At the outset we note that for all adaptation experiments we ran the limited memory BFGS algorithm for 100 iterations for each speaker. This number of iterations was determined in some preliminary experiments, and is sufficient to achieve convergence of

the WER. The secondary model is created by starting with the GMM corresponding to the full acoustic model and clustering (to minimize Kullback-Liebler divergence) to a desired number Gaussians.

We start our experiments when there are 4 Gaussians ($m = 4$) in the secondary model. Here, we initialize each \mathbf{A}_g , $g = 1, \dots, 4$, with an identity matrix of dimension 39 by 39. This configuration results a WER of 1.44%, which is worse than the standard FMLLR. We next initialize each \mathbf{A}_g by the standard FMLLR matrix before optimization. This gives a WER of 1.30% with $m = 4$.

Table 1 shows the performance when we use varying number of Gaussians to compute the posteriors. Again each matrix \mathbf{A}_g is initialized with the standard FMLLR matrix. We see that the best performance is obtained with about 8 Gaussians. As the number of Gaussians is reduced we have very few parameters and this hurts performance. As we increase the number of Gaussians beyond a certain point we expect to degrade because of over training. Clearly the optimal point will depend on the amount of data for a certain speaker. In our test set each speaker has the same amount of data so we did not experiment with varying the number of Gaussians per speaker.

As the next step we constraint the FMLLR matrices in order to reduce the number of parameters. Since the features are concatenation of basic 13 dimensional MFCC, delta and delta-delta features, we chose to constrain FMLLR matrices to block matrices. We optimize only the blocks of 13 by 13 dimensional sub-matrices on the diagonal for each \mathbf{A}_g . An example \mathbf{A}_g is given as

$$\mathbf{A}_g = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}$$

where \mathbf{I} is an identity matrix of size 13 by 13. With this constrain, the model with 16 Gaussians resulted a WER of 1.21%. We further increase the number of Gaussians to 64 and constrain each \mathbf{A}_g to a diagonal matrix. This configuration resulted a WER of 1.36%.

m	WER	Num. parameters
Baseline	2.08%	-
FMLLR	1.41 %	1571
4	1.30%	6084
8	1.25%	12168
16	1.32%	24336

Table 1. Adaptation with different number of Gaussians in secondary model. Each \mathbf{A}_g is initialized with the standard FMLLR matrix.

In calculating the posteriors in ϕ we could use an additional scale factor α as below:

$$\phi_g(\vec{x}_t) = \frac{\pi_g N(\vec{x}_t; \vec{\mu}_g, \Sigma_g)^\alpha}{\sum_k \pi_k N(\vec{x}_t; \vec{\mu}_k, \Sigma_k)^\alpha}.$$

We considered this option since choosing alpha appropriately can cause the posteriors to be smoother in their variation across time. Note that we could choose α by optimizing over α to maximize likelihood, which we did not do in this paper. The error rates at various α 's are shown in Table 2. All of these results use 8 Gaussians in the secondary model. As α goes to

α	WER
Baseline	2.08%
FMLLR	1.41%
0.02	1.29%
0.2	1.22%
1.0	1.25%
8.0	1.34%

Table 2. Adaptation with different scales in calculating secondary model posteriors

zero the performance degradation is expected since the posterior distribution is uniform in the limit. Although the total performance across scales is pretty much the same and close to optimal, we maybe able to further improve the performance by allowing the scale to be speaker dependent. Picking the best scale gives a total error rate of 1.22%. To practically obtain this improvement we could choose the scale to maximize likelihood, in fact we could let the scale be a parameter which is also optimized along with \mathbf{A} .

In our final set of experiments we tried to improve upon FMLLR with the non-linear adaptation technique introduced in this paper. Using just FMLLR we get to an error rate of 1.41%. We fix the FMLLR matrix \mathbf{A}_0 and then training the \mathbf{A} matrix to maximize likelihood. In this configuration the features used were: $\tilde{\mathbf{y}}_t = \mathbf{A} \tilde{\mathbf{f}}(\mathbf{A}_0 \tilde{\mathbf{x}}_t)$. Each \mathbf{A}_g is then initialized with an identity matrix. The results for varying number of Gaussians are listed in Table 3. We observe that fixing nor-

m	WER	Num. parameters
Baseline	2.08%	-
FMLLR	1.41 %	1571
4	1.30%	6084
8	1.25%	12168
16	1.28%	24336

Table 3. Adaptation with different number of Gaussians in secondary model. Adaptation after basic FMLLR. Each \mathbf{A}_g is initialized with the identity matrix.

malizing the features with FMLLR before the algorithm or initializing the algorithm with the FMLLR matrix gives similar results.

5. CONCLUSIONS AND FUTURE WORK

In this paper we introduced a non-linear adaptation technique based on the FMPE feature generation model [1]. We see

a 14% relative improvement using this non-linear technique over the standard FMLLR which is a 42% relative improvement over the baseline system. In the future we would like to explore the choice of $\tilde{\phi}$, use larger secondary GMMs in conjunction with a map to a smaller number of classes to constrain the number of parameters and to allow the parameters of the secondary model to be changed to maximize likelihood of the test data. We would also like to use this maximum likelihood approach during training, as opposed to the original FMPE paper which uses the MPE criterion.

6. REFERENCES

- [1] D. Povey, B. Kingsbury, L. Mangu, G. Saon, H. Soltau, and G. Zweig, "FMPE: discriminatively trained features for speech recognition," in *Proceedings of ICASSP*, 2005.
- [2] A. Acero and R. M. Stern, "Environmental robustness in automatic speech recognition," in *Proceedings of ICASSP*, 1990.
- [3] M. J. F. Gales, "Maximum likelihood linear transformations for HMM-based speech recognition," *Technical report, TR 291, Cambridge University*, 1997.
- [4] P. Olsen, S. Axelrod, K. Visweswariah, and R. A. Gopinath, "Gaussian mixture modeling with volume preserving non-linear feature space transforms," in *Proceedings of ASRU*, 2003.
- [5] S. Dharanipragada and M. Padmanabhan, "A nonlinear unsupervised technique for unsupervised adaptation," in *Proceedings of ICSLP*, 2000.
- [6] M. K. Omar and M. Hasegawa-Johnson, "Non-linear maximum likelihood transformation for speech recognition," in *Proceedings of Eurospeech*, 2003.
- [7] D. C. Liu, J. Nocedal, "On the limited memory BFGS method for large scale optimization problems," *Mathematical Programming*, vol. 45, pp. 503–528, 1989.
- [8] More, Thuente, "Line search algorithms with guaranteed sufficient decrease," *ACM TOMS*, vol. 20, no. 3, pp. 286–307, 1994.
- [9] M. S. Gockenbach, W. W. Symes, "The Hilbert Class Library," <http://www.trip.caam.rice.edu/txt/hcldoc/html/>.
- [10] S. Deligne, S. Dharanipragada, R. Gopinath, B. Maisson, P. Olsen, and H. Printz, "A robust high accuracy speech recognition system for mobile applications," *IEEE Transactions on Speech and Audio Processing*, pp. 551–561, 2002.