



CAMBRIDGE UNIVERSITY
ENGINEERING DEPARTMENT

**Automatic Transcription of
Conversational Telephone Speech
Development of the CU-HTK 2002 System**

Thomas Hain, Phil Woodland, Gunnar Evermann, Mark Gales, Andrew Liu, Gareth Moore, Dan Povey & Lan Wang

CUED/F-INFENG/TR.465

December 2003

Cambridge University Engineering Department
Trumpington Street
Cambridge. CB2 1PZ
England

E-mail: th223@eng.cam.ac.uk

<http://www-svr.eng.cam.ac.uk/~th223>

Abstract

This paper discusses the Cambridge University HTK (CU-HTK) system for the automatic transcription of conversational telephone speech. A detailed discussion of the most important techniques in front-end processing, acoustic modelling and model training, language and pronunciation modelling are presented. These include the use of conversation side based cepstral normalisation, vocal tract length normalisation, heteroscedastic linear discriminant analysis for feature projection, Minimum Phone Error Training and speaker adaptive training, lattice-based model adaptation, confusion network based decoding and confidence score estimation, pronunciation selection, language model interpolation and class based language models.

The transcription system developed for participation in the 2002 NIST Rich Transcription evaluations of English conversational telephone speech data is presented in detail. In this evaluation the CU-HTK system gave an overall word error rate of 23.9%, which was the best performance by a statistically significant margin. Further details on the derivation of faster systems with moderate performance degradation are discussed in the context of the 2002 CU-HTK 10×RT conversational speech transcription system.

1 Introduction

The transcription of conversational telephone speech is one of the most challenging tasks for speech recognition technology. State-of-the-art systems still yield high word error rates typically within a range of 20–30%. Work on this task has been aided by extensive data collection, namely the Switchboard-I corpus [10]. Originally designed as a resource to train and evaluate speaker identification systems, the corpus now serves as the primary source of data for work on automatic transcription of conversational telephone speech in English.

The first reported assessment of word recognition performance on the Switchboard-I corpus was presented in [9] with an absolute word error rate of around 78%¹. In this experiment only a small portion of the Switchboard-I corpus was used in training. Over the years the performance of systems on this task has gradually improved. Progress is assessed in the yearly “Hub5E” evaluations conducted by the U.S. National Institute for Standards in Technology (NIST). The Cambridge University HTK group first entered these evaluations in 1997 using speech recognition technology based on the Hidden Markov Model Toolkit (HTK) [37] and has participated in evaluations on this task ever since. This paper describes the CU-HTK system for participation in the 2002 NIST Rich Transcription (RT-02) evaluation. We focus on two test conditions: the unlimited compute transcription task where the only design objective is the word error rate (WER); and the less than 10 times real-time (10×RT) transcription task where the system processing time is not allowed to exceed 10 times the duration of the speech signal.

The paper is organised as follows: the first section briefly reviews basic aspects of the HTK Large Vocabulary Recognition (LVR) system, followed by a detailed description of the data used in experiments. In Section 4 we present the acoustic modelling techniques essential to our system and discuss particular data modelling aspects. Section 5 outlines the pronunciation modelling, followed in Section 6 by a description of the language models used in our systems. In Section 7 we discuss issues in decoding and system combination. The structure of the full transcription system is presented in Section 8, including a detailed analysis of the performance on large development and evaluation test sets. This system served as the basis for the 10×RT system described in Section 9.

2 HTK LVR Systems

The HTK large vocabulary speech recognition systems are built using the Hidden Markov Model Toolkit [37] and are based on context dependent state clustered HMM sets with Gaussian mixture output distributions. The same basic model training methodology is used for a variety of tasks. The acoustic data is normally represented by a stream of 39 dimensional feature vectors with a frame spacing of 10ms, based on 12 MF-PLP coefficients [33] and the zeroth cepstral coefficient c_0 representing the signal energy. The first and second order derivatives of each coefficient are appended to form the full feature vector. The words are mapped into phoneme strings using dictionaries based on a modified and regularly version of the LIMSI 1993 WSJ pronunciation dictionary [8]. The dictionaries contain multiple pronunciations per word. Cross-word context-dependent phone models using a context of either ± 1 in the case of triphones or ± 2 for the quinphones are used as the acoustic models. In addition to models for speech, the acoustic model set usually contains two silence models, one for silence, and one for short inter-word pauses with the latter preserving context across words. In order to avoid under-training and the effect of unseen phone contexts, the HMM states are clustered using phonetic decision trees trained with a maximum likelihood criterion [37]. Initial single Gaussian per state models are created prior to state clustering by two-model re-estimation [35]. After state clustering and several iterations of Baum-Welch re-estimation, the number of mixture components is gradually increased, interleaved with multiple re-estimation steps.

The language models (LMs) are based on N-grams with backoff for smoothing. For conversational telephone speech large amounts of task-dependent material is not available and other

¹The focus of the work was topic and speaker identification rather than word recognition.

schemes have to be adopted (see Section 6).

In order to allow for the use of complex acoustic and language models word graphs (lattices) are used extensively in re-scoring passes. The recognition system normally operates in multiple stages where an important aspect is acoustic model parameter adaptation, for example using maximum likelihood linear regression [21]. A more detailed description of the basic setup can be found in [35] or [37].

3 Training and Evaluation Data

The experiments in this paper made use of data from the *Switchboard-1* (Swbd1) corpus [10], the CallHome English corpus and small parts from the Switchboard-2 corpus. The Switchboard-1 corpus covers more than 2900 conversations from about 550 U.S. speakers. The speakers were unknown to each other and were requested to converse on a certain topic. The data used had 4-wire recordings with a sample rate of 8kHz and μ -law encoded with a resolution of 8 bits per sample. An initial manual segmentation and transcription of more than 250 hours of speech was provided by the Linguistic Data Consortium (LDC). The inherent difficulty even in manual transcription of this type of data had forced many research sites to create their own segmentations and corrected versions of the transcriptions. In order to provide an improved baseline the data was more recently re-transcribed and re-segmented by Mississippi State University (MSU)². The *CallHome English* (CHE) corpus consists of a total of 200 conversations between family members and close friends, no restriction was placed on the conversation topics. The LDC distributed 120 conversations, comprising a total of 18.5 hours of speech. The remainder was used in Hub5E evaluations in the years 1997-2001. A particular, though minor effect is the occasional existence of multiple speakers per conversation side on this data³. The *Switchboard-2* (Swbd2) corpus was collected with the intention to serve as test-bed for speaker recognition, consequently most of it is not transcribed. The corpus was collected in a total of 5 phases, all calls were made within the U.S.A. The phases were recorded in different regions of the U.S.A, the fourth phase is also called Switchboard Cellular (Cell) collecting data over mobile phone channels with special focus on the GSM channel⁴. Each corpus has specific unique attributes and ASR system performance varies significantly with the corpus from which the data is drawn. In the following word error results are also presented for each of the data sources in the test sets.

Multiple training and test sets are used in the experiments in this paper. The selection of data for training of acoustic models allows a scaling of the complexity of experiments. Table 1 shows details of the training sets used in this paper. Note that most experiments are based on the `h5train02` set which covers data from Swbd1, CHE and Cell.

The Switchboard-1 part of `h5train00` and consequently `h5train00sub` and `h5train02` are based on a January 2000 snapshot of the MSU transcripts. The segment boundaries for these sets have been modified to reduce the amount of silence present in the training data. Based on forced alignments of the training data, a collar of only 0.2 seconds of silence on either side was retained and segments were split at long pauses. The table only shows data used for acoustic training. Details on the data used for language model training data can be found in Section 6.

Table 1 also shows the test-sets used in this paper. Note that `dev01` is the official 2001 Hub5 evaluation development set [28] consisting of 40 sides of Switchboard-2 (from the 1998 evaluation), 40 sides of Switchboard-1 (from the 2000 evaluation) and 38 sides of Switchboard-2 cellular data. The `dev01sub` set was selected to show similar word error rates to the full `dev01` set. For all cases a manual segmentation into speaker turns was available.

²See <http://www.isip.msstate.edu/projects/switchboard/index.html>

³A total of 12 conversation sides have multiple speakers with a maximum of 3 speakers on one side

⁴A more detailed description of the data can be obtained from the LDC web-site: <http://www ldc.upenn.edu>.

Table 1: Data sets used for training and test. Training sets are described by the source of transcripts, test sets by the originating corpus.

	Dataset	Description	#hours
Train	<code>h5train98</code>	LDC(Swbd1,CHE)	180
	<code>h5train00</code>	MSU(Swbd1), LDC(CHE)	265
	<code>h5train00sub</code>	Subset of <code>h5train00</code>	68
	<code>h5train02</code>	<code>h5train00</code> + LDC(Cell)	282
Test	<code>eval98</code>	Swbd1/CHE	3
	<code>dev01</code>	Swbd1/Swbd2/Cell	6
	<code>dev01sub</code>	Half of <code>dev01</code>	3
	<code>eval02</code>	Swbd1/Swbd2/Cell	6.5

4 Acoustic modelling experiments

In this section fundamental acoustic modelling techniques for conversational telephone data are presented. We discuss front-ends, the use of feature transformation schemes, data issues, the use of discriminative and speaker adaptive training schemes, and test-set speaker adaptation.

4.1 Acoustic Analysis

Due to the special transfer characteristics of telephone channels, the lower and upper frequency regions of the speech signal are attenuated and often very greatly so. In order to avoid the placement of filter-banks in regions containing only noise, the frequency analysis has been restricted to a range of 125-3800Hz [13]. Initial experiments indicate WER improvements using these band limits after cepstral normalisation.

4.1.1 Cepstral mean and variance normalisation

Cepstral mean normalisation (CMN) can be used to reduce the effects of constant channel characteristics. The normalisation can be performed on a per segment basis if the segments are of sufficient length. However, the audio data for this task has an average utterance length of 3.4 seconds (`h5train02`) which also includes a collar of 0.2 seconds of silence at the boundaries. Therefore segment based mean normalisation is problematic. Since acoustic conditions can be assumed to be relatively stable over the duration of a conversation the mean can be calculated over a complete conversation side. This approach will be referred to as side-based mean normalisation (Side-CMN). In preliminary experiments [12] significant reductions in WER by about 1% absolute with Side-CMN compared to segment-based CMN were observed.

In a similar manner to CMN, variance normalisation can also be used. Again normalisation on a per-side basis (Side-CVN) is advisable. Initial results indicate a 1-1.5% absolute improvement with both Side-CVN and Side-CMN over Side-CMN only. Another important advantage of side-based CVN with respect to its effect on vocal tract length normalisation are discussed in the following section. For more detailed results the reader is referred to [12, 13].

4.1.2 Vocal Tract Length Normalisation

Maximum likelihood vocal tract length normalisation (VTLN) implements a per speaker frequency scaling of the speech spectrum [20]. The optimal scale factor α_{opt} is obtained by searching for the factor that yields the highest data likelihood. The optimal scale factor is then applied to yield a speaker specific feature stream. Normalisation can be performed on the test data only or both on the training and test data. The advantage of VTLN lies in its simplicity and effectiveness. Since




Figure 1: Piecewise linear frequency scaling. Two parameters are needed to determine the frequency scaling are the scale factor α , and the cut-off frequency f_U .

VTLN	dev01
—	42.1
test	41.1
train & test	38.0

Table 2: %WER for systems trained on `h5train98` and tested on `dev01`. VTLN warping in test only or training and test.

only a single parameter needs to be estimated, the method is robust and can be implemented very efficiently.

The frequency scaling can be implemented by inverse scaling of the Mel filter-bank centre frequencies. In [12] we proposed a piecewise linear approach of the form presented in Fig. 1. This form ensures that frequencies tie up at the boundaries. The cut-off frequency is determined in advance. Warp factors are found by searching over a range of warp factors where the data likelihood is computed by performing an alignment of a previously obtained word level transcript. In our experience the quality of this transcript only has a minor effect on performance given enough speaker-specific data is available. Note that in this process the cepstral mean and variance normalisation vectors have to be recomputed. Warping with a certain factor has an effect on the data likelihoods that would introduce a systematic bias in the warp factor search. This bias should be corrected using Jacobian compensation, but the application of CVN to the warped data achieves the same effect.

Table 2 shows a performance comparison of VTLN in test only and in both training and test. Triphone acoustic models and a trigram language model were used in the experiments. The gain from test-only VTLN is less than half the gain obtained when VTLN is used in both training and test. Overall a relative reduction in WER of about 10% is usually observed over a range of test sets. In order to obtain a reasonable warp factor distribution on the training data multiple iterations of warp factor estimation and model retraining are necessary. Fig. 2 shows a typical distribution of warp factors across speakers on the `h5train02` training set. Note the clear split in the warp factor per gender with a broad distribution within-gender.

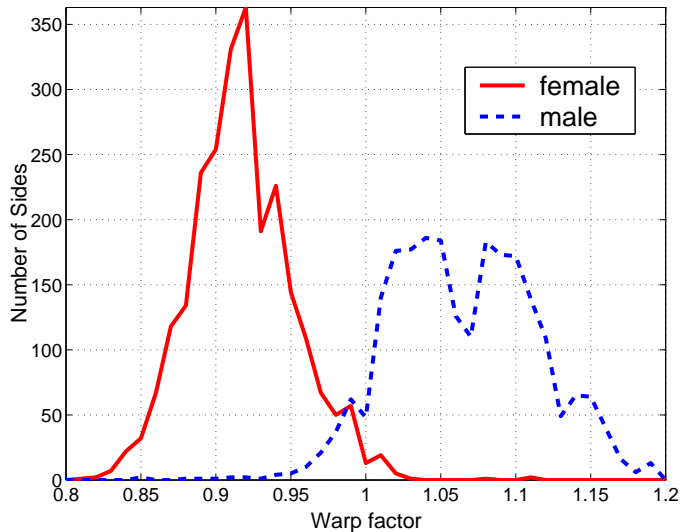


Figure 2: Frequencies of warp factors per gender on `h5train02`. Triphone models for likelihood measurement were trained on `h5train98`.

Table 3: %WER on `dev01sub` using 16comp VTLN ML triphones. Results are obtained using re-scoring of 4-gram lattices.

Data	Swbd1	Swbd2	Cellular	Total
<code>h5train00</code>	25.2	42.1	42.5	36.5
<code>h5train02</code>	24.9	41.3	41.7	35.8
<code>h5train02w</code>	24.9	41.0	41.4	35.7

4.2 Acoustic training data

The training sets used for transcription of conversational telephone speech corpus are relatively large compared to those available on other corpora. Such large training sets are required due to the considerable amount of variability in the data. In [14] we showed that, not unexpectedly, the incremental gain from an increase in the amount of training data slowly decreases even on a logarithmic scale. Starting from 20 hours of data trebling the amount resulted in a WER improvement of more than 4%, a further trebling of the amount of data only gave an additional 1.6% gain⁵.

Another aspect is the appropriateness for the acoustic conditions. Table 3 shows experimental results on `dev01sub` using triphone models trained on different training sets. Note that the `dev01sub` test set also contains data from Swbd2 and Cell. The WER on the cellular data is similar to the Swbd2 performance. By adding the about 17 hours of cellular data to the training set (`h5train02`) the word error rate can be reduced by 0.7% absolute. Note that the improvement is mostly on the Swbd2 and Cell portions of the data.

The amount of training data available from the Cellular and CHE corpora is small. Hence a second system was built with the Cellular and CHE data weighted by a factor of 2 (`h5train02w`). A slight reduction in WER by 0.1% was observed with more substantial improvement on the Cell part of the test set. For all `h5train02` ML trained systems in this paper the data weighting was

⁵It is important to note that these results were obtained by testing on Swbd2 and CHE data. The smallest training set however only contained data from Swbd1, while the larger sets included CHE data

Table 4: %WER on dev01sub using 28comp h5train02 triphones, with MF-PLP (Std), semi-tied (ST) covariance matrix and HLDA front-ends. Results were obtained by re-scoring of 4-gram lattices.

	Feature Transform	Swbd1	Swbd2	Cellular	Total
ML	—	24.2	40.8	40.4	35.1
	ST	23.0	39.2	39.3	33.7
	HLDA	22.2	38.8	39.1	33.3
MPE	—	21.0	37.0	36.6	31.4
	HLDA	19.3	35.6	35.8	30.1

maintained.

All experimental results in Table 3 were obtained using triphone models with 16 mixture components per speech state and approximately 6000 decision tree clustered states. Additional experiments with higher numbers of components gave an optimum result at 28 mixture components. This further decreased the overall error rate to 35.1%.

4.3 Heteroscedastic LDA

In this work a mixture of Gaussians with diagonal covariance is used to model the output distribution of each state. However, it is known that there are correlations in the feature vector which may limit the ability of the mixture model to accurately represent the underlying data. One solution to this is to linearly transform the features so that the data associated with each component is approximately uncorrelated. In addition, some dimensions contain little discriminatory information and should be removed.

Heteroscedastic linear discriminant analysis (HLDA) [18] is a linear projection scheme and may be viewed as a generalisation of LDA. It removes the restriction that all the within class covariance matrices are the same. The HLDA projection matrix, \mathbf{A} , for a d -dimensional feature space, \mathbf{o}_t , may be written as

$$\hat{\mathbf{o}}_t = \mathbf{A}\mathbf{o}_t = \begin{bmatrix} \mathbf{A}_{[p]}\mathbf{o}_t \\ \mathbf{A}_{[d-p]}\mathbf{o}_t \end{bmatrix} \quad (1)$$

where the top p dimensions are deemed to be those dimensions that contain discriminatory information, the useful dimensions. The final $(d-p)$ -dimensions, contain no useful information and are called the nuisance dimensions. Those are modelled using a global distribution and hence can be ignored for the purpose of further training and evaluation.

The maximum likelihood estimate of the transform parameters can be obtained in an iterative process [7]. In this work the projection matrix is initialised to an identity matrix. The useful dimensions were selected based on Fisher ratios.

The HLDA transforms built for this work projected a 52-dimensional feature down to 39 dimensions.

The 52 dimensional source vector consisted of the standard 39-dimensional feature vector with third order derivatives appended. HMMs with 16-component mixtures based on the standard feature vectors were trained and then extended to incorporate the third derivatives. After further steps of Baum-Welch re-estimation the transform was estimated and the number of model components increased by mixing-up [37] to 28 components. For semi-tied covariance systems [7], the process is identical except no addition of third derivatives was used.

Table 4 compares systems using the 39-dimensional front-end with the use of both a global semi-tied covariance system, and HLDA system and with ML training. Using a global semi-tied covariance matrix reduced the error rate by 1.4% absolute. An additional 0.4% absolute was obtained by using HLDA with third order differential coefficients rather than a semi-tied

system. Hence there is additional discriminatory information that can be extracted from the third derivatives. Discriminative training (see Section 4.4) was then applied to the ML-trained system. Though the reduction in WER due to HLDA with discriminatively trained models is 1.3% absolute rather than the 1.8% obtained with the ML-trained system, there was still a significant advantage in using HLDA.

4.4 Discriminative training

The standard criterion for estimation of HMM parameters is maximum likelihood (ML). The maximum likelihood estimate is optimal in the sense that it is consistent with minimum variance. Two important assumptions are made: a large number of training samples is available; and the model itself is correct, i.e. reflects the true nature of the data. Neither of these assumptions is true for HMM based speech recognisers [26]. Consequently discriminative training schemes are of interest.

In conjunction with HMM based speech recognisers several discriminative training schemes have been proposed. Most importantly we have shown that Maximum Mutual Information (MMI) yields better performance than ML for the transcription of conversational telephone speech. Use of discriminative criteria in training is well known in small vocabulary tasks [29]. The next section gives a brief description of MMI training, followed by a more detailed description of an alternative discriminative training scheme, Minimum Phone Error (MPE) training.

4.4.1 Maximum Mutual Information

For R training observation sequences $\{\mathcal{O}_1, \dots, \mathcal{O}_r, \dots, \mathcal{O}_R\}$ with corresponding transcriptions $\{s_r\}$, the MMI objective function for HMM parameter set λ , including the effect of scaling the acoustic and LM probabilities can be written

$$\mathcal{F}_{\text{MMI}}(\lambda) = \sum_{r=1}^R \log \frac{p_{\lambda}(\mathcal{O}_r | \mathcal{M}_{s_r})^{\kappa} P(s_r)}{\sum_s p_{\lambda}(\mathcal{O}_r | \mathcal{M}_s)^{\kappa} P(s)} \quad (2)$$

where \mathcal{M}_s is the composite model corresponding to the word sequence s , $P(s)$ is the probability of this sequence as determined by the language model and κ is a scale factor⁶. The summation in the denominator of (2) is taken over all possible word sequences allowed in the task. Hence MMI maximises the posterior probability of the correct sentences. The denominator in (2) can be approximated by a word lattice of alternative sentence hypotheses.

The Extended Baum-Welch (EBW) algorithm is used for parameter optimisation [11]. The parameter update formulae require the collection of numerator (num) and denominator (den) statistics derived from word lattices based on recognition of the training set. In order to allow a broader range of confusable hypotheses the use of a weak language model has been shown to be important [34]. The parameter update formulae for the means and variances of Gaussian m in state j is given by

$$\begin{aligned} \hat{\mu}_{jm} &= \frac{\{\theta_{jm}^{\text{num}}(\mathcal{O}) - \theta_{jm}^{\text{den}}(\mathcal{O})\} + D_{jm}\mu_{jm}}{\{\gamma_{jm}^{\text{num}} - \gamma_{jm}^{\text{den}}\} + D_{jm}} \\ \hat{\sigma}_{jm}^2 &= \frac{\{\theta_{jm}^{\text{num}}(\mathcal{O}^2) - \theta_{jm}^{\text{den}}(\mathcal{O}^2)\} + D_{jm}(\sigma_{jm}^2 + \mu_{jm}^2)}{\{\gamma_{jm}^{\text{num}} - \gamma_{jm}^{\text{den}}\} + D_{jm}} - \hat{\mu}_{jm}^2 \end{aligned} \quad (3)$$

$\theta^{\text{num}}(\mathcal{O}^n)$ and $\theta^{\text{den}}(\mathcal{O}^n)$ denote n -th order Gaussian occupancy weighted sums of the data based on the numerator and denominator lattices respectively and γ is the Gaussian occupancy summed over time. The constant D can be used to control convergence speed and robustness and is set on a per Gaussian level.

⁶It is assumed that the LM probabilities $P(s)$ have already been ‘‘scaled’’ (raised to the power) by the normal LM scale factor $1/\kappa$ and hence further scaling by κ takes them back to their original values.

Table 5: %WER on `eval98` and a subset of the training set (`train`) using triphone models. `train` results were obtained using a lattice unigram LM, test-set results by re-scoring of 4-gram lattices.

Criterion	Training set	τ^I	<code>train</code>	<code>eval98</code>
ML	<code>h5train00sub</code>	—	47.8	46.5
MMI	<code>h5train00sub</code>	50	32.2	43.8
ML	<code>h5train00</code>	—	47.2	45.6
MMI	<code>h5train00</code>	200	35.8	41.4
MPE	<code>h5train00</code>	100	34.4	40.8

It was shown that data weighted per Gaussian interpolation between ML and MMI parameter estimates substantially increases the robustness of the training process. I-smoothing [31] is a way of applying an interpolation between a ML and a discriminative objective function in a way which depends on the amount of data available for each Gaussian. In the case of MMI this means that the numerator occupancies are increased by a certain amount τ^I , while leaving the average first and second order data values unchanged. I-smoothing is used in training of all discriminative models in this paper.

Table 5 shows a comparison of ML versus MMI, on different training set sizes. Note that the ML system serves as a starting point for further parameter estimation steps using the MMI criterion. In conjunction with I-smoothing substantial improvements in WER are obtained. I-smoothing trades performance on the training set against improved generalisation to the test set. Furthermore the relative gain from using MMI training on a 68 hour training set is about 6%, however on 265 hours of data the relative WER improvement is 10%.

4.4.2 Minimum Phone Error Training

The aim in discriminative training is to choose the model parameters such as to minimise the word error rate. The Minimum Word Error Rate (MWE) criterion [31] is designed to maximise the expected word accuracy on the training set. It was found that an equivalent formulation at the phone level yields better generalisation. The Minimum Phone Error criterion is defined as

$$\mathcal{F}_{\text{MPE}}(\lambda) = \sum_{r=1}^R \frac{\sum_s p_\lambda(\mathcal{O}_r | \mathcal{M}_s)^\kappa P(s) \text{RawAccuracy}(s)}{\sum_s p_\lambda(\mathcal{O}_r | \mathcal{M}_s)^\kappa P(s)} \quad (4)$$

where $\text{RawAccuracy}(s)$ is a measure of the number of phones accurately transcribed in hypothesis s . The objective function takes into account that many phone sequences are at least partially correct and consequently cannot fully count as a competitor. This is expressed by adding the correct fraction to the numerator. Details on how the computation of the $\text{RawAccuracy}(s)$ and the re-estimation procedure can be found in [31].

Table 5 also shows WER results on `eval98` for MPE training using the `h5train00` training set. MPE outperforms MMI with I-smoothing by 0.6% absolute on this test set. Improvements of a similar kind have been verified on other test sets. It is important to note the necessity for I-smoothing when using the MPE criterion. Without I-smoothing the absolute WER is 50.7% which is even poorer than the baseline ML result. More detailed results and descriptions of MPE/MWE are given in [31] and [36].

4.5 Unsupervised test-set adaptation

Adaptation to the test speaker and the acoustic environment greatly improves the performance of automatic speech recognisers. The VTLN and side based normalisation schemes discussed above are essentially adaptation techniques and show substantial performance improvements. Techniques

Table 6: %WER on dev01sub using 28comp triphones trained on h5train02. All results are obtained by rescoring of 4-gram lattices.

	Adaptation	Swbd1	Swbd2	Cellular	Total
ML	-	22.2	38.8	39.1	33.3
	global MLLR	20.2	35.8	36.4	30.7
MPE	-	19.3	35.6	35.8	30.1
	global MLLR	18.0	33.6	34.3	28.5

discussed in this section use errorful transcripts of the data generated in previous decoding stages for adaptation supervision.

4.5.1 Maximum Likelihood Linear Regression

MLLR [21] is a well known adaptation technique that is widely used in many speech recognition tasks. The model parameters are adjusted by a linear transform (including a bias) such that the likelihood on the adaptation data is increased. On this task, assuming a single speaker on a conversation side, side-based adaptation is performed. Typically mean vectors are adapted using block-diagonal transforms. However, the use of HLDA (see Section 4.3) removes the association of blocks to the static cepstra and the derivatives. Consequently a full transform for the means was used in these experiments. Variance adaptation is performed using diagonal transforms [5].

The sets of Gaussians can be grouped into multiple classes, each associated with a particular transform. This is especially helpful if more data is available. The classes can be found using regression class trees [22] or by manual definition. In practise most of the WER gain is obtained by using one transform for speech and one for the silence models, further termed global MLLR.

Table 6 shows results using a single iteration of global MLLR using ML and MPE trained model sets. Note that the models also use side-based mean, variance and VTL-normalised data. The relative improvement using ML trained models is about 8%, compared to a 5% relative reduction in WER using MPE trained models. The relative improvements are similar for all sub-sets of the data.

4.5.2 Full variance transforms

In the above experiments diagonal transforms for variance adaptation were used. Even with multiple transform classes this guarantees that the model covariances remain diagonal after adaptation. Improved performance can be obtained by using a full matrix for speaker based variance adaptation. In this case the transformed covariance matrix, $\hat{\Sigma}^{(m)}$, for mixture component m is given by

$$\hat{\Sigma}^{(m)} = \mathbf{H}\Sigma^{(m)}\mathbf{H}^\top \quad (5)$$

where $\Sigma^{(m)}$ is the speaker-independent covariance matrix and \mathbf{H} is a linear transform estimated on the adaptation data. This can be interpreted as speaker-dependent global semi-tied covariance matrix [7]. A summary of this adaptation scheme is given in [6]. There are a number of options. Theoretically the transforms may be full, diagonal, or block-diagonal. In practice a full transform was used in all cases in this paper. The full variance (FV) transform was computed after standard mean and variance MLLR. Typically a WER reduction of 0.4% to 0.8% was obtained. However as a side effect, we found that there were reduced benefits from multiple MLLR regression classes when used with a full variance transform.

4.5.3 Lattice based MLLR

In unsupervised adaptation the word level output of a previous recognition pass is used as the word level supervision. Errors in this transcription will affect the effectiveness of the adaptation.

Table 7: %WER on dev01sub using 28-component HLDA MPE triphones trained on h5train02. All results are obtained by rescoring of 4-gram lattices.

HLDA	Adapt	Swbd1	Swbd2	Cellular	Total
×	-	21.0	37.0	36.6	31.4
×	LatMLLR+FV	18.6	33.7	34.6	28.9
✓	-	19.3	35.6	35.8	30.1
✓	MLLR	18.0	33.6	34.3	28.5
✓	LatMLLR	17.5	32.7	32.9	27.6
✓	LatMLLR+FV	17.6	32.4	32.7	27.5

To compensate for the uncertainty in the transcription a set of alternative word hypotheses can be used as the supervision. The lattice MLLR technique presented in [32] employs word lattices to represent these alternatives. During the transform estimation the contribution of the alternatives is weighted by posterior probability estimates based on a lattice forward-backward pass.

Table 7 shows a break-down of results for the adaptation techniques on dev01sub. When using HLDA, global MLLR adaptation brings an improvement of 1.6% WER absolute over the baseline. Iterative lattice MLLR using 2 speech transforms brings a further 0.9%. In this case the result can be improved only slightly when using a FV transform. Overall the improvement in WER from adaptation in this case is about 9% relative. Table 7 also shows results of adaptation with or without the use of an HLDA transform (Section 4.3). Note that the difference between systems without or with HLDA transforms for unadapted models is 1.3% WER absolute, compared to 1.4% with adapted models. This indicates that the improvements from these techniques are approximately additive.

4.6 Speaker adaptive training

Adaptive training is a powerful training technique for building speech recognition systems on non-homogeneous data. Variability in the training data may result from the speaker changing, differing acoustic environments or varying channel conditions. The basic idea of adaptive training is to use one or more transformations of features or model parameters to represent these speaker and environment differences. A canonical model can then be trained, given the set of speaker/environment transforms. This canonical model should be more compact and amenable to being transformed to a new speaker, or acoustic condition, than standard speaker independent (SI) systems.

Forms of adaptive training have already been incorporated into the training process in the form of VTLN and mean and variance normalisation (see Section 4.1). These methods use constrained transformations of the feature space for normalisation. However, gains are increased further by incorporating linear transformations of model parameters, for example MLLR [21], into the training process. This was the original form of speaker adaptive training (SAT) described in [1]. One of the issues with the original SAT scheme is the cost of training. This may be solved by using constrained MLLR transforms [6]. Then SAT can be implemented by transforming the features, and there is no need to change the model parameter optimisation process. As constrained MLLR is a feature space transform it is simple to incorporate it into the discriminative training framework.

Table 8 shows the performance of SAT systems, trained using constrained MLLR transforms. In testing, the systems are adapted by using constrained MLLR. An improvement of 0.5% absolute over the baseline can be observed, for the ML trained models, where the gain originates from the more difficult data Switchboard-2 and Switchboard Cellular. With MPE training the gain is reduced. Note that these results, for the purpose of consistency, only involve test set adaptation using constrained MLLR. Using the MPE HLDA SAT models as in Table 8 in conjunction with LatMLLR+FV adaptation gives a WER of 27.3% absolute. The use of SAT is important to yield complimentary system output for the purpose of system combination(see Section 7.2).

Table 8: %WER on dev01sub using 28comp HLDA triphones trained on h5train02. All results are obtained by rescoring of 4-gram lattices using constrained MLLR adaptation only.

	SAT	Swbd1	Swbd2	Cellular	Total
ML		20.5	36.0	36.4	30.9
	✓	20.5	35.1	35.9	30.4
MPE		17.9	33.4	34.0	28.3
	✓	18.1	33.3	33.6	28.2

5 Pronunciation modelling

CU-HTK systems use a pronunciation dictionary for translation of words into phoneme sequences, where each word in the dictionary has one more more possible pronunciations associated with it. The dictionaries used in training and test are obtained from the CU base dictionary. The core of this base is the 1993 LIMS WSJ lexicon [8], with manually generated additions and corrections. Pronunciations for new words if needed are added manually. On average 1.1 to 1.2 pronunciations per word are included.

5.1 Pronunciation Probabilities

Unigram pronunciation probabilities, that is the probability of a certain pronunciation variant for a particular word, can be estimated based on an Viterbi-alignment of the training data. Counting the number of occurrence of pronunciation variants gives rise to an estimate for the probabilities. Considerable smoothing is necessary to account for the inevitable data sparsity.

The dictionaries in the HTK system explicitly contain silence models as part of a pronunciation. This allows the inclusion of the silence models when estimating probabilities. The most successful scheme in our experiments uses three separate dictionary entries for each real pronunciation which differed by the word-end silence type: no silence; a short pause preserving cross-word context; and a general silence model altering context. Smoothing of the probability estimates used the overall distribution for each silence variant. Finally all dictionary probabilities are renormalised so that for a given word the probability of the most likely pronunciation is set to one to avoid an additional penalty for words with many variants. During recognition the pronunciation probabilities are scaled by the same factor as used for the language model. Table 9 shows that the use of pronunciation probabilities gives a reduction in WER of 0.9% absolute on dev01sub for ML. Even larger gains were observed on some other test-sets [15].

5.2 Single Pronunciation (SPron) Dictionaries

The standard approach to pronunciation modelling is to use multiple pronunciations (MPron) for each word. However, the considerable pronunciation variation in conversational data makes the use and selection of multiple pronunciations difficult and causes additional confusability in decoding. Theoretically Gaussian mixture based HMMs should be able to cope with phone or sub-phone substitutional effects. These phone substitutions are the main cause of the existence of multiple pronunciations in dictionaries. In this case the training of model parameters can implicitly perform a similar task to manual phonemic labelling.

An automated scheme for deriving a single pronunciation from the multiple pronunciation dictionary was developed. This is described in detail in [16]. The algorithm obtains pronunciation information from the acoustic training data to train simple statistical models that allow the selection of pronunciation variants. Since the list of words used in training usually differs from that used in recognition, the algorithm also provides a method for the selection of pronunciations

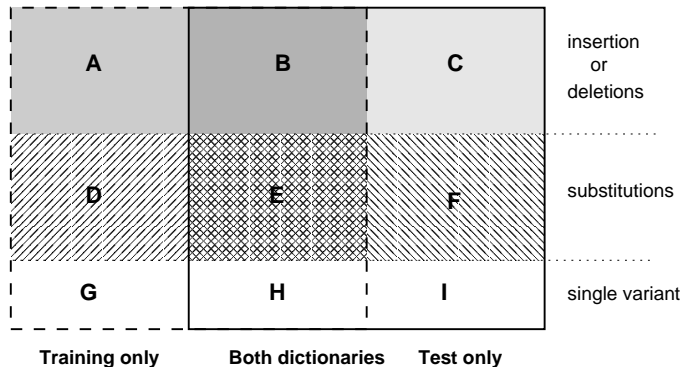


Figure 3: A joint view of the list of words used in training and test.

Table 9: %WER on dev01sub using 28comp triphone models trained on h5train02, HLDA and optionally pronunciation probabilities. Results are obtained by rescoring of 4-gram lattices.

	Training	PrProb	Swbd1	Swbd2	Cell	Total
MPron	ML	×	22.2	38.8	39.1	33.3
	MPE	√	21.5	37.9	38.1	32.4
		×	19.3	35.6	35.8	30.1
		√	19.1	35.0	35.6	29.8
SPron	ML	×	21.6	37.9	37.8	32.3
	MPE	√	21.3	37.7	37.4	32.0
		×	19.4	35.2	35.1	29.8
		√	19.6	34.9	34.9	29.7

for words not observed in training. An MPron dictionary and an HMM set trained using that dictionary are used to obtain pronunciation frequencies from the training data.

The words in the training and test vocabularies can be partitioned into 9 categories as shown in Figure 3. Words with a single pronunciation are left unchanged (G, H and I). For words observed in the training data (A, B, D and E) unseen pronunciation variants are removed and variants are sorted according to frequency. For these words variants with phoneme substitutions only (D and E) are merged. Finally, the most frequent variant is retained, in the case of equal frequency a random choice is made. For words not observed (C and F) the observed words (A, B, D and E) are used to form a statistical model. This model allows variants with phoneme substitutions to be merged and then those variants predicted to be less frequent by the model can be deleted. The precise form of the model used is given in [16]. Using this single pronunciation dictionary the system is rebuilt, including the generation of the phonetic decision tree.

Table 9 shows a comparison of the SPron system with the standard MPron system, and in addition the use of pronunciation probabilities. Note that for the SPron system the “pronunciation probability” is simply the probability of the word being followed by an optionally deletable “short” silence model, or a standard silence model. In the ML training case the SPron system outperforms the baseline MPron system by 1% absolute, in conjunction with pronunciation probabilities this is reduced to 0.4% absolute. In the case of MPE training both pronunciation probabilities and SPron system give reduced gains. However, the output of SPron and MPron systems still differ significantly and consequently can be used for system combination (see Section 7.2).

Table 10: Perplexities on various test sets using interpolated language models (tg=trigram,fg=4-gram).

Model	N-gram	eval98	dev01cell	eval02
AcAllLM	fg	90.9	78.4	77.4
CellLM	fg	112.9	91.5	104.1
BNLM	fg	111.6	106.5	121.0
fgint02	fg	72.4	63.8	64.1
ClassLM	tg	112.9	99.7	102.6
fgintcat02	fg	71.5	62.8	63.3

6 Language Modelling and Word Lists

In most speech transcription tasks such as for example dictation or Broadcast News transcription large amounts of representative text data are available. In the case of transcription of spontaneous speech over the telephone a very large corpus is not available as the cost of transcribing the data is considerable. Consequently the amount of in-domain data for the training of language models and vocabulary selection is fairly small and restricted to the transcription of the acoustic training data.

CU-HTK systems have followed a certain strategy for building language models and selecting word lists on this task for several years [13]. All words from the acoustic training set are used in decoding⁷. In order to minimise the Out-of-Vocabulary (OOV) rate on the test sets, this set of words is merged with the 50000 most frequent words occurring in 204 million words (MW) of Broadcast News (BN) training data, yielding a vocabulary size of around 55000. Given the vocabulary, word bigram, trigram and 4-gram language models are trained on the acoustic LM training set. These models are then interpolated with corresponding models trained on the BN corpus. The resulting 4-gram LM is further interpolated with a class-based trigram language model where the classes are automatically generated based on word bigram statistics [17, 24, 27].

The vocabulary selected for the CU-HTK 2002 Hub5E system is derived from `h5train02` in the aforementioned manner, giving a word list of 55449 distinct entries. This yielded an OOV rate of 0.38% on `eval98`, 0.17% on the cellular part of `dev01` and 0.54% on `eval02`. Three language models were trained on different sets of data: `AcAllLM` is trained on `h5train00` and `h5train98` to encompass both MSU and LDC transcript styles. This version of the transcripts includes the false starts and covers a total of 6.2MW. As the amount of training was relatively small the model was trained using modified Kneser-Ney discounting [2]. `CellLM` was trained on the Switchboard Cellular part of the `h5train02` set with a size 0.2MW. Again modified Kneser-Ney discounting was used. The `BNLM` model was trained on 204MW of Broadcast News data ranging in epoch from January 1992 to December 1997 to cover approximately the dates of data collection. Smoothing used Katz-backoff and Good-Turing discounting. The individual language models were merged to form a single language model that effectively interpolates the component models with interpolation weights 0.43:0.25:0.32 for the three language models (`AcAllLM:CellLM:BNLM`) where interpolation weights were chosen by perplexity minimisation. The merged language model (`fgint02`) contained 4.77 million bigrams, 6.33 million trigrams and 7.35 million 4-grams.

The class trigram language model used 350 automatically generated word classes. Classes and trigram models were trained on the `h5train02` transcriptions only. The final class model contained 75k bigrams and 337k trigrams. In the final interpolation stage the optimal weight was 0.81 for the word 4-gram and 0.19 for the class model. Table 10 shows the perplexities on several test sets. Note that before merging the individual perplexities are high compared to the merged model (`fgint02`). Despite significantly higher perplexities the category LM yields a further reduction of

⁷Note that for the purpose of LM training some text processing steps are necessary to deal with for example partial words.

about one point in perplexity on all test sets.

7 Decoding and System Combination

In Section 2 we briefly mentioned the generic decoding strategy for HTK LVR systems. Initial Viterbi decoding passes are used to produce lattices that allow the search space to be constrained in subsequent passes. This is useful for using more complex acoustic and language models, for example using quinphone models or 4-gram LMs.

7.1 Minimum word error rate decoding

The standard criterion in ASR uses the Maximum A Posteriori (MAP) principle. For continuous speech recognition this implies the search for the sentence yielding the highest posterior probability. This is notably different from the desired objective of word error rate (rather than sentence error rate) minimisation. The use of confusion networks [23] allows an efficient implementation of the minimum word error rate principle. For a particular word lattice link posterior probabilities are estimated using the forward-backward algorithm. The lattice is transformed into a linear graph, or confusion network (CN) employing a link clustering procedure [23]. Both temporal information as well as information on phonetic similarity of words is used in clustering. The linear graph consists of a sequence of so called confusion sets, which contain competing single word hypotheses with associated posterior probabilities. By picking the word with the highest posterior from each set the sentence hypothesis with the lowest overall expected word error rate can be found. The use of CN decoding normally reduces the WER by 0.5%-1% absolute. More detailed results on the CU-HTK 2002 Hub5E evaluation system can be found in Section 8.2.

The estimates of the word posterior probabilities encoded in the confusion networks can be used directly as word confidence scores. As the posteriors tend to be over-estimates of the true posteriors they are mapped to confidence scores using a piece-wise linear function based on a decision tree [3].

7.2 System combination

In recent years interest in the development of complementary systems, i.e. systems that substantially differ in their word level output while retaining a similar word error rate, was stimulated by techniques such as ROVER [4]. ROVER allows the combination of the system outputs either by voting or by the use of confidence scores. This approach can be generalised to the use of confusion networks in the combination [3]. In this case CN output from each system is generated and dynamic programming is used to align the confusion networks. The cost function is an estimate for the probability of a word match given two confusion sets. After alignment the networks are merged and standard CN-decoding is applied. Confusion network combination (CNC) allows the weighting of systems, however normally with limited effect. In this paper no weighting is used in system combination. Results of CNC are discussed in Section 8.2.

8 The CU-HTK April 2002 Hub5E System

In the previous sections we have presented a set of techniques that are important for the transcription of conversational telephone speech and have discussed the performance of each technique. However, in practise the performance improvements are rarely additive and the selection of techniques is non-trivial. Accordingly the use of new techniques cannot be assessed purely using baseline comparisons, their operation in a complete speech recognition system is of at least equal importance. Consequently research of the development of large ASR systems is not only interesting for the purpose of finding the best performance, but also for an improved understanding of the relationship of techniques.

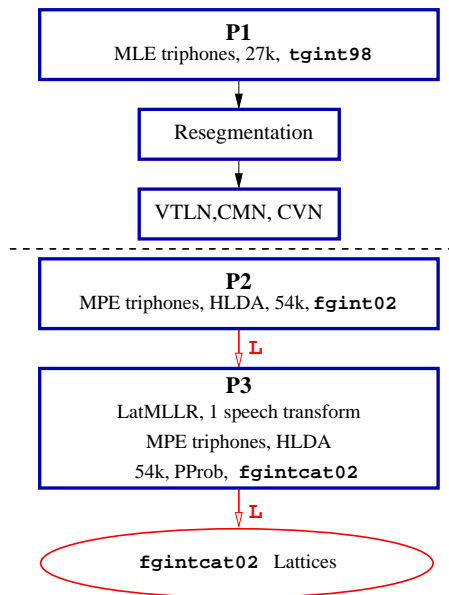


Figure 4: The first three passes (P1) - (P3) in the CU-HTK 2002 Hub5E system. Normal arrows denote word level output, arrows with “L “ denote lattices. The output of these stages is a set of word lattices.

The following describes the system designed for participation in the RT-02 evaluations for transcription of conversational telephone speech⁸. We discuss the broad structure and the processing and performance of the system.

8.1 System Structure

The system operates in a total of 6 passes. In general each pass employs techniques that have not been used in previous passes. A pass can sometimes be split into a certain number of sub-processes. The output of a pass is either a word string, a word lattice or a confusion network. Input to the system is a set of audio files and time information of segments to be recognised. The system presented here is based on manual segmentation of the audio stream. Each stage of the system operates on a conversation side basis⁹. Thus the following discussion will describe the transcription of a single conversation side only.

The first part of the system is shown in Fig. 4 and is designed for the generation of lattices to be used in the second part. In the first step the aim is to perform a robust estimation of mean and variance normalisation vectors and VTLN warp factors. The data is encoded in MF-PLP coefficients as described in Section 4.1. A set of transcripts is generated with ML triphone models and a trigram language model (for details see [13]). As this part of the system originates from the 1998 CU-HTK system, the word error rate of this pass is fairly high. However, no word level information is passed on to later stages. The initial set of word level transcripts also allows a re-segmentation of the data¹⁰. CMN and CVN vectors are recomputed and the new segments and the word level transcripts are used in VTLN estimation as described in 4.1.2. The estimated warp

⁸The interested reader is referred to <http://www.nist.gov/speech/tests/rt/rt2002>.

⁹This means that no information from the second channel or other conversations is used in the transcription process.

¹⁰Only 0.2 seconds of silence are left at segment boundaries and segments are split if more than a second of silence occurs.

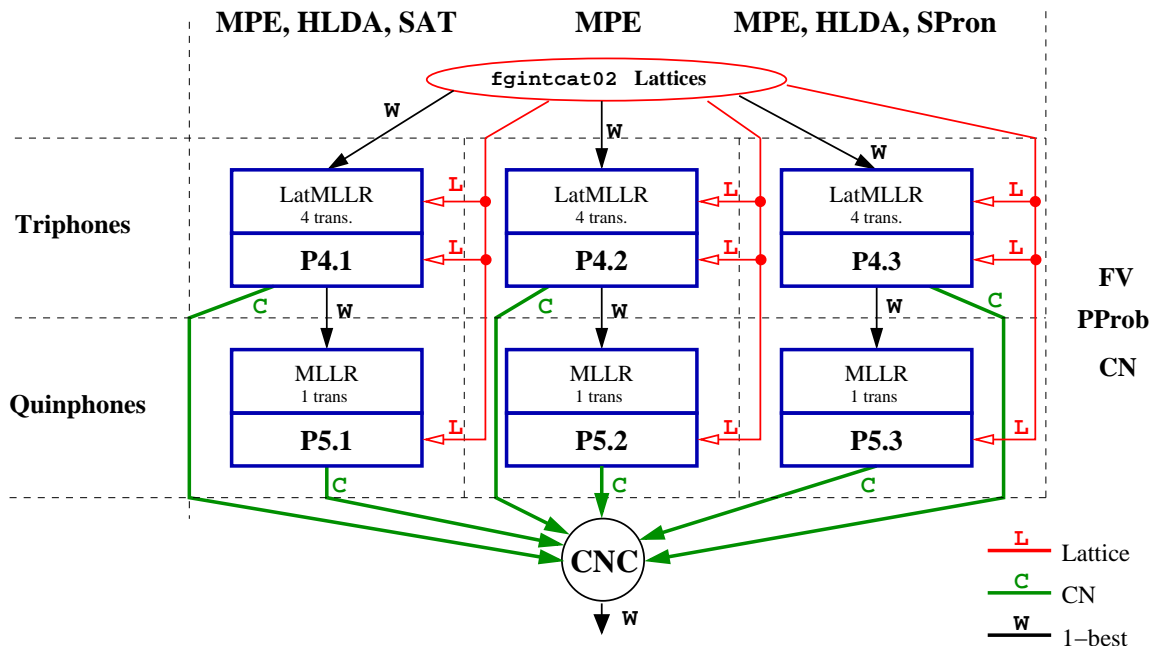


Figure 5: Stages P4. [123], P5. [123] and the system combination stage P6 of the CU-HTK 2002 Hub5E system.

factor is used to recode the audio data¹¹. Note that the re-segmented audio data is only used for estimation purposes, further steps use the original segmentation.

In the second pass (P2) transcripts for the purpose of HMM parameter adaptation are produced. For this purpose state clustered triphone models with 6155 speech states are used. Initially the models were trained on `h5train02` using the ML criterion using the standard HTK mix-up procedure [37]. After estimation of an HLDA matrix, the number of mixture components was increased to the total number of 28 Gaussians per speech state. Further re-estimation steps using the MPE criterion with I-smoothing gave the model set used in this pass. Decoding used the 54k dictionary described in Section 6. The Viterbi decoding pass used a trigram language model (`tgint02`) produced in the fashion described in Section 6. Further rescoreing using the `fgint02` LM gave the output of this pass.

In the third pass (P3) a set of lattices for use in all subsequent rescoreing passes is generated. Global MLLR transforms are estimated using the output from P2, followed by lattice-based MLLR. The acoustic models used are identical to those in P2 and the interpolated and class smoothed 4-gram LM `fgintcat02` was used to obtain a set of lattices. In addition pronunciation probabilities were used in decoding. Fig. 4 summarises the essential components of the first three stages. The output of P3 forms the basis for all subsequent processing steps.

The second part of the system is based on the principle of system combination and contains stages based on complementary systems that differ in model training and adaptation strategies. This part is split into branches where each branch corresponds to a specific model construction strategy. In total three branches were used: one associated with MPE SAT-trained models, together with HLDA feature projection and a standard MPron dictionary (branch 1); one associated with a non-HLDA MPE trained model (branch 2); and one branch using models based on an HLDA transform, MPE training and an SPron dictionary (branch 3). All branches are further subdivided into two passes: The (P4. [123]) passes are based on triphone models, the (P5. [123]) passes use quinphone models. Fig. 5 shows an outline of the essential operation blocks and the data flow.

¹¹If multiple speakers are present on a conversation side this procedure leads to undesired results. Nevertheless this effect is hardly encountered in the corpora used.

For adaptation in the triphone stage all branches use the same strategy: the `fgintcat02` lattices obtained in the first part of the system and the associated first best output is used in a lattice-based MLLR+FV scheme as described in Section 4.5.3. A total of four transforms for the speech classes are estimated. The decoding step of these branches consists of an acoustic rescoring step using the `fgintcat02` lattices generated in (P3) with the adapted models. Confusion networks were generated and word strings obtained using CN decoding. The acoustic models in the first two branches have the same characteristics in terms of the number of states and mixture components as those used in (P2)/(P3). The models in the third branch differ in so far as the number of system parameters is slightly smaller with a total of 5955 states with 28 mixture components each.

The quinphone models use a cross-word context of ± 2 phones with word boundary information and are trained in a similar fashion to the triphones. The average number of states associated with speech models is higher, with 9640 states for the models used in the first two branches and 9418 states in case of the SPron based models. Models are first trained up to 16 mixture components. After estimation of HLDA matrices (branches 1 and 3 models only) the number of mixture components was increased to 24. In the (P5. [123]) stages adaptation is performed using global MLLR, together with a full-variance transform. As the use of full cross-word quinphones substantially increases the size of static phone networks the quinphone decoding stages use the dynamic network decoder [30] for the rescoring of lattices.

The output of each of the stages (P4. [123]) and (P5. [123]) is a set of confusion networks. These are merged in the final confusion network combination stage (P6). In this stage minimum word error rate decoding is used to arrive at the final word level transcription for the conversation side. The overall structure of this second system part is represented in Fig. 5. Note that the arrows denote the flow of information, either in the form of word strings (W), lattices (L), or confusion networks (C). The final output is a word string with word level times and confidence scores.

8.2 Recognition Performance

Table 11 shows WERs on the full `dev01` test set and the full 2002 evaluation set for all system passes. Since performance of the individual stages is very similar for both test sets the following discussion will be concentrated on results on `dev01`. As `dev01` has served in development this shows that the results generalise to independent test sets.

The WER of the first pass (P1) is rather high. As no word level information is transferred beyond this stage this is of little concern. Re-segmentation removed a total of 2628 seconds of audio or about 12% of the audio data to obtain consistent amounts of silence at segment boundaries. The second pass (P2) shows the unadapted performance of triphone models using VTLN, HLDA, MPE training and a 4-gram language model. Note that the word error rate on Switchboard-1 data is about 20%, substantially lower than the data originating from Switchboard-2 and Cellular sources. The (P2) output lattices are used to estimate a global MLLR transform in (P3). Initial lattices are produced using the bigram language model. In subsequent lattice expansion and pruning steps more complex language models are applied. A more than 3% absolute performance difference can be observed between bigram and trigram language models. The use of 4-grams yields another 0.7% absolute. Smoothing with the class-based LM gave only a slight improvement, mostly due to performance on the Cell data. On the evaluation set performance gains are similar apart from applying the class LM.

As discussed above the second part of the system is split into three branches. The lowest WER is obtained in branch 1 using the MPE-SAT-HLDA model sets. Compared to (P3) the gain from using SAT, lattice MLLR with multiple speech transforms and FV transforms is 0.7% WER absolute. A further WER reduction by 0.7% is obtained when using CN decoding. For contrast purposes, the results for acoustic models as used in (P3) but with (P4. [123]) style adaptation was included in the table (labelled ‘‘Contrast’’). Note that the effective gain from SAT after CN decoding is 0.3% absolute. The output of stage (P4. 1) is used for adaptation in (P5. 1) which, after CN decoding, gives an absolute WER of 25.7%. The second branch yields poorer performance due to lower complexity of the model set, but was found to be useful in system combination. Branch 3 differs from the first by the use of a SPron dictionary and non-SAT training, obtaining the same

Table 11: %WER on dev01 using h5train02 training data for all stages of the CU-HTK 2002 Hub5E System.

Passes	Comment	CN	dev01			eval02				
			Swbd1	Swbd2	Cell	Total	Swbd1	Swbd2	Cell	Total
Part 1	P1	×	31.7	46.9	48.1	42.1	35.6	44.6	50.5	44.0
	P2	×	20.1	34.7	34.3	29.6	24.6	30.9	34.8	30.4
	P3	×	22.0	36.2	35.1	31.0	26.3	31.5	35.4	31.4
Part 2	P4.1	×	18.0	31.4	30.2	26.5	22.1	26.7	30.6	26.8
	P4.2	✓	17.5	30.7	29.6	25.9	21.6	26.3	29.6	26.1
	P4.3	×	19.8	32.7	32.4	28.2	23.7	28.5	32.8	28.6
Contrast	tri	×	18.4	31.5	30.0	26.5	22.3	27.4	31.2	27.2
	tri	✓	18.0	31.0	29.7	26.2	22.1	27.3	29.7	26.6
	tri	×	18.2	32.0	30.7	26.8	22.0	27.0	30.6	26.8
P5.1	SAT quin	×	17.9	31.6	30.1	26.4	22.5	26.3	29.9	26.5
	SAT quin	✓	17.2	30.8	29.2	25.7	21.5	25.5	28.6	25.4
	non-HLDA quin	×	19.5	32.7	32.0	28.0	23.7	27.7	32.3	28.2
P5.2	non-HLDA quin	✓	18.5	31.8	30.6	26.9	22.4	26.7	30.7	26.9
	SPron quin	×	18.6	31.8	29.7	26.6	22.2	27.0	29.8	26.6
	SPron quin	✓	18.1	31.1	28.8	25.9	21.5	26.4	28.8	25.8
CNC	quin	×	18.1	32.0	30.3	26.7	22.3	26.6	30.0	26.6
	quin	✓	17.5	31.2	29.2	25.9	21.4	25.8	28.8	25.5
	P4. [123] + P5. [123]	✓	16.4	29.2	27.4	24.2	19.8	24.3	27.0	23.9

Table 12: Execution times of decoding stages using IBM x330 servers (with Pentium III 1GHz processors)

Pass	P1	P2	P3	P4. [123]	P5. [123]
Speed (\times RT)	12	11	37	131	147

results as the contrast system. Note that in general the error rates on Cell data are lower for this branch. The gain from CN decoding is on average 0.7% absolute for the first branch, 1.1% on the second and 0.3%-0.7% for the third. The quinphone stages give only marginal improvements over the results of the triphone stages. Compared to the contrast system the performance of the SAT models is slightly better whereas the SPron quinphones give identical word error rates. The value of these systems lies in their contribution to system combination. Combining the output of the triphone stages (P4. [123]) gives a WER of 24.9% whereas the final result of CNC of all system output is 24.2%, or a 1.5% absolute gain from system combination.

The performance of the individual passes on the evaluation set is similar. The performance of the SPron triphone models was better giving the lowest triphone word error rate. Overall the gain from CN-decoding of quinphone model output was higher, especially in the case of SAT models with 1.1% WER absolute. The final word error rate of 23.9% was the lowest in the NIST RT-02 evaluations by a statistically significant margin [19]. The confidence scores obtained from confusion networks gave a normalised cross-entropy value (see e.g. [28]) of 0.289 on *eval02*.

Table 12 shows the execution times for the recognition stages associated with each of the passes. The individual numbers exclude times for estimation of lattice-based MLLR. The overall system had a real-time (RT) factor of about 320. In comparison the result on *eval02* after CN-decoding of (P3) output lattices is 26.7% WER absolute using only $67 \times$ RT.

9 Building Faster Systems

The system presented in the previous section was designed for optimal WER performance. In practical scenarios it is not feasible to take minutes of processing time for the transcription of a second of speech. Consequently there is considerable interest in research on how to modify recognition systems to yield optimal performance under computational constraints. One test-condition in the 2002 NIST RT Evaluation focused on operation in less than $10 \times$ RT. The system described before was modified to meet this objective. The first part of the full system has relatively low complexity. Thus this part was chosen to form the basis of development. Several issues were important in development: the processing stages with relatively low gain but high computational costs were excluded (for example lattice-based MLLR); Lattices allow fast decoding with complex models or adaptation strategies, but a three stage approach to lattice generation is too costly; pruning parameters in the decoding stages can be tuned to substantially decrease the real-time factor with moderate degradation in WER performance; the use of faster computers with local disk storage allows for considerable speed improvements.

The final system is structured as follows: The first stage is identical to the full system P1 pass, however, much tighter pruning in decoding is used. In the second stage fast adaptation using least squares linear regression (LSLR) is performed. Using the HLDA MPE triphone models lattices are produced with the interpolated trigram LM *tgint02* and further re-scored using *fgint02*; the output of this stage is used as supervision in adaptation using standard MLLR with two transforms for speech models and one for silence. Lattices are re-scored and CN decoding is performed on the output lattices.

Table 13 shows WER results for the CU-HTK RT-02 $10 \times$ RT evaluation system. The high error rate of the first pass gives only poor supervision, a second MLLR based rescoreing step allows further improvements. Note that the final result of 27.2% is only 0.5% absolute from the CN

Table 13: %WERs and real time factors of the CU-HTK 2002 10×RT system on the eval02 test set. Speed as measured using an AMD Athlon XP 1900+.

Stage	Comment	%WER	Speed (×RT)
P1-10x	initial trans., VTLN	45.2	1.65
P2-10x	LSLR, lat-gen, fgint02	28.5	5.36
P3-10x	MLLR, re-score, CN	27.2	2.24

output of the full system (P3) stage which took 67×RT.

10 Conclusions

We have presented a complete state-of-the-art system for the transcription of conversational telephone speech and we described a range of techniques in acoustic, pronunciation and language modelling specifically important for this task. Particularly powerful methods in acoustic modelling are the use of side-based cepstral normalisation, VTLN, discriminative training using the MMI or MPE criteria, and heteroscedastic LDA. Speaker adaptation using standard or lattice-based MLLR and full variance transforms yields considerable word error rate improvements. In language modelling the use of a background Broadcast News corpus together with class based language models allows to compensate for the lack of training data on this task. Pronunciation probabilities give consistent performance improvements. The use of lattices allow the use of confusion network decoding and the efficient implementation of system combination. We have discussed several systems with similar performance and their use in system combination.

Overall the word error rate achievable on the original Swbd1 corpus is now below 20%. More natural data is available in the form Swbd2 and Cell data where the error rates are just below 30% absolute. The reasons for these high error rates are manifold and can only partly be attributed to lack of data. Undoubtedly error rates are still too high for many applications but development of speech recognition systems for this task is an ongoing process. Recently another Rich Transcription evaluation was held in March 2003, partly using new and more challenging data.

Acknowledgements

Many people have contributed to the development of HTK LVR systems in our group and thus have indirectly helped in putting this system together. The authors would like to especially acknowledge the work of Thomas Niesler and Ed Whittaker who helped to build earlier Hub5 systems. The work on the 2002 system was in part supported by GCHQ. Part of the experimental work for this paper was supported by DARPA grant MDA972-02-1-0013. The paper does not necessarily reflect the position or the policy of the US Government and no official endorsement should be inferred.

References

- [1] T. Anastasakos, J. McDonough, R. Schwartz, & J. Makhoul (1996). A compact model for speaker-adaptive training. In *Proc. ICSLP*, pp. 1137–1140.
- [2] S.F. Chen & J. Goodman (1998). An Empirical Study of Smoothing Techniques for Language Modeling. *Tech-report TR-10-98*, Computer Science Group, Harvard University.
- [3] G. Evermann & P.C. Woodland (2000). Posterior Probability Decoding, Confidence Estimation and System Combination, In *Proc. Speech Transcription Workshop*, College Park, MD.

- [4] J.G. Fiscus (1997). A Post-Processing System to Yield Reduced Word Error Rates: Recogniser Output Voting Error Reduction (ROVER). *Proc. IEEE ASRU Workshop*, pp. 347–354, Santa Barbara.
- [5] M.J.F. Gales & P.C. Woodland (1996). Mean and Variance Adaptation within the MLLR Framework. *Computer Speech & Language*, Vol. 10, pp. 249–264.
- [6] M.J.F. Gales (1998). Maximum Likelihood Linear Transformations for HMM-Based Speech Recognition. *Computer Speech & Language*, Vol 12, pp. 75–98.
- [7] M.J.F. Gales (1999). Semi-tied covariance matrices for hidden Markov models. In *IEEE Transactions Speech and Audio Processing*, 7:272–281.
- [8] J.-L. Gauvain, L.F. Lamel, G. Adda & M. Adda-Decker (1994). The LIMSI Nov93 WSJ System. In *Proc. SLT'94*, pp125-128, Plainsboro, NJ.
- [9] L. Gillick, J. Baker, J. Baker, J. Bridle, M. Hunt, Y. Ito, S. Lowe, J. Orloff, B. Peskin, R. Roth, & F. Scattone (1993), Application of Large Vocabulary Continuous Speech Recognition to Topic and Speaker Identification using Telephone Speech. In *Proc. ICASSP'93*, pp. 471 - 474.
- [10] J.J. Godfrey, E.C. Holliman, J. McDaniel (1992). SWITCHBOARD: Telephone Speech Corpus for Research and Development, In *Proc. ICASSP'92*, pp. 517–520.
- [11] P.S. Gopalakrishnan, D. Kanevsky, A. Nadas, & D. Nahamoo (1991). An Inequality for Rational Functions with Applications to Some Statistical Estimation Problems. In *Trans. Information Theory*. Vol 37, pp. 107-113.
- [12] T. Hain & P. C. Woodland (1998). CU-HTK Acoustic modelling experiments. In *Proc. NIST Hub5 Workshop*, Linthicum Heights, MD.
- [13] T. Hain, P.C. Woodland, T.R. Niesler, & E.W.D. Whittaker (1998). The 1998 HTK system for transcription of conversational telephone speech. In *Proc. ICASSP'99*, pp. 57–60.
- [14] T. Hain, P.C. Woodland, G. Evermann & D. Povey (2000). The CU-HTK March 2000 Hub5E Transcription System. *Proc. Speech Transcription Workshop*, College Park, MD.
- [15] T. Hain, P.C. Woodland, G. Evermann & D. Povey (1999). New Features in the CU-HTK System for Transcription of Conversational Telephone Speech.. *Proc. ICASSP'01*, Salt Lake City.
- [16] T. Hain (2003). Implicit Modelling of Pronunciation Variation in Automatic Speech Recognition. Submitted to *Speech Communication*.
- [17] R. Kneser & H. Ney (1993). Improved Clustering Techniques for Class-Based Statistical Language Modelling. In *Proc. Eurospeech'93*, pp. 973-976, Berlin.
- [18] N. Kumar (1997). Investigation of Silicon-Auditory Models and Generalization of Linear Discriminant Analysis for Improved Speech Recognition, Ph.D. Thesis, Johns Hopkins University, Baltimore.
- [19] A. Le & A. Martin (2002). The 2002 NIST RT Evaluation Speech-To-Text Results. Rich Transcription Workshop 2002, <http://www.nist.gov/speech/tests/rt/rt2002/presentations/index.htm>.
- [20] L. Lee & R.C. Rose (1996). Speaker Normalization using Efficient Frequency Warping Procedures. In *ICASSP'96*, pp. 353-356.

- [21] C.J. Leggetter & P.C. Woodland (1995). Maximum Likelihood Linear Regression for Speaker Adaptation of Continuous Density HMMs. *Computer, Speech and Language*, Vol. 9, pp. 171–186.
- [22] C.J. Leggetter & P.C. Woodland (1995), Flexible Speaker Adaptation using Maximum Likelihood Linear Regression. In *Proc. Eurospeech'95*, pp. 1155–1158, Madrid.
- [23] L. Mangu, E. Brill & A. Stolcke (1999). Finding Consensus Among Words: Lattice-Based Word Error Minimization. In *Proc. Eurospeech'99*, pp. 495–498, Budapest.
- [24] S. Martin, J. Liermann & H. Ney (1995). Algorithms for Bigram and Trigram Clustering. In *Proc. Eurospeech'95*, pp. 1253–1256, Madrid.
- [25] J. McDonough & W. Byrne (1999), Single-pass Adapted Training with All-pass Transforms. In *Proc. EuroSpeech'99*, pp. 2737–2740, Budapest.
- [26] A. Nadas, D. Nahamoo & M.A. Picheny (1988), On a Model-Robust Training Algorithm for Speech Recognition. In *IEEE Trans. ASSP*, Vol. 36, pp. 1432–1435.
- [27] T.R. Niesler, E.W.D. Whittaker & P.C. Woodland (1998). Comparison of Part-Of-Speech and Automatically Derived Category-Based Language Models for Speech Recognition. In *Proc. ICASSP'98*, pp. 177–180, Seattle.
- [28] The NIST Speech Group (2000), The 2001 NIST Evaluation Plan for Recognition of Conversational Speech over the Telephone. Version 1.1, www.nist.gov/speech/tests/ctr/h5.2001/h5-01v1.1.pdf.
- [29] Y. Normandin (1991). An Improved MMIE Training Algorithm for Speaker Independent, Small Vocabulary, Continuous Speech Recognition. In *Proc. ICASSP'91*, pp. 537–540, Toronto.
- [30] J.J. Odell, V. Valtchev, P.C. Woodland & S.J. Young (1994). A One Pass Decoder Design For Large Vocabulary Recognition. In *Proc. 1994 ARPA Human Language Technology Workshop*, pp. 405–410, Morgan Kaufmann.
- [31] D. Povey & P.C. Woodland (2002), Minimum Phone Error and I-Smoothing for Improved Discriminative Training, In *Proc. ICASSP'02*, Orlando.
- [32] L.F. Uebel & P. C. Woodland (2001). Speaker Adaptation Using Lattice-based MLLR. In *Proc. ISCA ITRW Adaptation Methods for Speech Recognition*, Sophia-Antopolis.
- [33] P.C. Woodland, M.J.F. Gales, D. Pye & S.J. Young (1997). Broadcast News Transcription using HTK. In *Proc. ICASSP'97*, pp. 719–722, Munich.
- [34] P.C. Woodland & D. Povey (2000). Large Scale Discriminative Training for Speech Recognition. In *Proc. ISCA ITRW ASR2000*, pp. 7–16, Paris.
- [35] P.C. Woodland (2002), The development of the HTK Broadcast News transcription system: An overview. In *Speech Communication*, Vol. 37, pp. 47–67.
- [36] P.C. Woodland, G. Evermann, M.J.F. Gales, T. Hain, X.L. Liu, G.L. Moore, D. Povey, & L. Wang (2002), CU-HTK April 2002 Switchboard System, In *Proc. Rich Transcription Workshop*, Vienna, VA.
- [37] S.J. Young, G. Evermann, T. Hain, D. Kershaw, G.L. Moore, J.J. Odell, D. Ollason, D. Povey, V. Valtchev, P.C. Woodland (2003). The HTK Book. Cambridge University, <http://htk.eng.cam.ac.uk>.