The 2006 RWTH Parliamentary Speeches Transcription System

J. Lööf, M. Bisani, Ch. Gollan, G. Heigold, B. Hoffmeister, Ch. Plahl, R. Schlüter, and H. Ney

Lehrstuhl für Informatik 6 - Computer Science Dept.
RWTH Aachen University, Aachen, Germany
{loof,bisani,gollan,heigold,hoffmeister,plahl,schlueter,ney}@cs.rwth-aachen.de

Abstract

In this work, investigations in the course of the developement of RWTH automatic speech recognition systems developed for the second TC-STAR evaluation campaign 2006 are presented. The systems were designed to transcribe parliamentary speeches taken from the European Parliament Plenary Sessions (EPPS) in European English and Spanish, as well as speeches from the Spanish Parliament. The RWTH systems apply a two pass search strategy with a fourgram one-pass decoder including a fast vocal tract length normalization variant as first pass. The systems further include several adaptation and normalization methods, minimum classification error trained models, and bayes risk minimization. For all relevant individual components contrastive results are presented on the EPPS Spanish and English data, including investigations which did not yet enter the evaluation systems.

1 Introduction

The TC-STAR (Technology and Corpora for Speech to Speech Translation) project (TcS, httpwwwtc starcom) is envisioned as a long-term effort to advance research in all core technologies for Speech-to-Speech Translation (SST), including Automatic Speech Recognition (ASR), Spoken Language Translation (SLT) and Text to Speech (TTS) (speech synthesis). The project targets a selection of unconstrained conversational speech domains (speeches and broadcast news) and three languages (British English, European Spanish, and Mandarin Chinese). For the TC-STAR project, language resources (LR) for English and Spanish parliamentary speeches were collected for training and system development, as well as the TC-STAR evaluation campaigns. Within the restricted conditions of the TC-STAR evaluations the training data is restricted to these LR. This paper describes in detail the English and Spanish RWTH ASR system which were developed for the restricted condition of the TC-STAR Second Evaluation Campaign 2006. The systems comprises a one-pass fourgram decoder including fast vocal tract length normalization (VTLN), constrained maximum likelihood linear regression (CMLLR) including speaker adaptive training (SAT), maximum likelihood linear regression (MLLR), discriminative training including minimum classification error (MCE) training, as well as *Bayes* risk minimization (MBR). Furthermore, internal system combination, including ROVER, confusion network combination (CNC), as well as a new lattice-based system combination procedure were investigated and compared. Also open vocabulary recognition was considered. Nevertheless, for the evaluation tasks both internal system combination as well as open vocabulary recognition were not yet considered, since preliminary experiments did not yet show consistent improvements.

2 Language Resources

2.1 Data

The English and Spanish LR both contain recordings from the European Parliament Plenary Sessions (EPPS), whereas the Spanish LR additionally include speeches from the Spanish Parliament and Congress (SPC). Approximately 100h of speech recordings per language were manually transcribed. These verbatim transcriptions (VT) include a segmentation into sentence like units, speaker labels, and topic headings. Although most of the speeches are planned, almost all speakers exhibit the usual effects known from spontaneous speech, like hesitations, false starts and articulatory noises. These disfluencies are also annotated.

The web site of the European Parliament (EUR, httpwwweuroparleuint) provides all EPPS reports since April 1996 translated in all official languages of the EU. These documents are known as the final text edition (FTE) and differ notably from the VT as the FTE aims for high readability.

Table 1 specifies the data used for language modelling. The recordings, the corresponding manual transcripts, and the text LR were produced by Universitat Politècnica de Catalunya (UPC) and RWTH Aachen University. Table 2 gives the statistics of the acoustic training data used in the RWTH system. Development and evaluation sets were provided by ELDA. The English and Spanish EPPS development and evaluation sets each consisted of about three hours of speech, plus 4h of Spanish parliament data for evaluation. Table 3 gives an overview of the corpora.

Table 1: *Text resources available for language modelling.*

	running words		
	transcriptions	FTE	Spanish SPC
English	781,649	33,894,405	-
Spanish	516,936	35,190,383	47,181,386

Table 2: *Transcribed recordings from the EPPS (both) and SPC (Spanish) domain available for acoustic modelling.*

	English	Spanish
Acoustic Data [h]	87.5	91.3
# Segments	66,670	101,608
# Running Words	704,883	743734

2.2 Lexicon Modeling

The recognition word lists were derived from the restricted domain data as described in Sec. 2.1. The available textual data was cleaned up and normalized, using a manually defined set of rules and semi-automatic methods. The word

Table 3: *Development and evaluation data from the EPPS domain, and from the SPC domain (for evaluation Spanish only).*

	English		Spanish	
	Dev	Eval	Dev	Eval (+SPC)
Audio [h]	3.2	3.2	2.4	6.9
# Run. wrd	27,029	29,829	20,982	60,039
# Speakers	41	41	31	63
Vocab size	52,429		60,156	
4-gram PP	99.7	108.7	78.2	88.9
OOV [%]	0.81	0.58	0.61	1.22

lists were produced as follows. All words from the verbatim transcriptions occuring at least twice were chosen. For the additional textual data a cut-off value was calculated requiring an out-of-vocabulary (OOV) rate below one percent on the development and a final lexicon of at least 50k words.

The English pronunciation lexicon was derived from the British English Example Pronunciation Dictionary (BEEP). The Spanish pronunciation lexicon was derived from the lexicon of the LC-STAR project (LcS, httpwwwlc starcom). Using the dictionaries statistical grapheme-to-phoneme conversion models were trained (Bisani and Ney, 2003) for Spanish and English. The models were used to produce pronunciations for words not covered by the original lexica. In Table 3 the lexicon statistics on the development and evaluation data are presented.

2.3 Handling of OOV Words

We had originally assumed that the EPPS task would exhibit a very high lexical diversity leading to inevitably high out-of-vocabulary (OOV) rates. To address this problem an open vocabulary recognition approach was examined. In this so-called flat hybrid approach we augment the recognition vocabulary by a set of word fragments each consisting of a short sequence of phonemes with associated spelling information. The set of fragments is derived from the baseline pronunciation dictionary using a maximum likelihood criterion. The language model used by the recognizer is estimated from a modified version of the training corpus where each OOV word is replaced by its most likely sequence of fragments. This technique has been applied quite successfully on the ”Wall Street Journal” database (Bisani and Ney, 2005). We have tried the identical techniques on the EPPS data, however without success: OOV words were recognised only very rarely, while spurious insertions of small fragments increased the overall error rate. We attribute this failure to the surprisingly low lexical diversity of the EPPS task. A low frequency of OOV words in training causes the estimation of the ”OOV part” of the model to be unreliable, due to lack of data. At the same time a low OOV rate in testing means that false alarms may easily exceed the potential improvement from OOV detection.

3 Acoustic Modeling

3.1 Baseline Acoustic Modeling

The acoustic front end comprises Mel-Frequency Cepstral Coefficient (MFCC) features derived from a bank of 20 filters. 16 cepstral coeficients including the zeroth coefficient were used, and cepstral mean normalization was applied. The MFCC features were augmented with a *voicing feature* (Zolnay et al., 2002). The MFCCs and voicing features from nine consecutive frames were concatenated and a linear discriminative analysis (LDA) was used to project the resulting vector to 45 dimensions.

Acoustic models were triphone based Gaussian mixture models (GMMs) with a globally pooled diagonal covariance matrix. The triphones were top down clustered using CART, rendering 4501 generalized triphone states.

The baseline acoustic models were maximum likelihood (ML)/*Viterbi* trained using the manually transcribed training data provided for the restricted condition, cf. Sec. 2.1 and Table 2.

3.2 Speaker Normalization and Adaptation

Three different approaches were used in combination to compensate for the acoustical variations due to speaker differences. First, a fast one-pass variant of Vocal Tract Length Normalization (VTLN) (Eide and Gish, 1996)(Welling et al., 1999) was applied to the filterbank within the MFCC extraction both in training and testing. The fast VTLN performs warping factor estimation using GMMs trained on a subset of the training corpus, for which warping factors were estimated using the usual grid search. Speaker adaptive training (SAT) based on Constrained Maximum Likelihood Linear Regression (CMLLR) (Gales, 1998) was used to compensate for speaker variation in both training and testing. The *Simple Target Model* (STM) approach (G. Stemmer, 2005) was used, since results in (G. Stemmer, 2005) indicate that it outperforms the standard CMLLR-SAT method (Gales, 1998). As target model an acoustic model with a single Gaussian per state trained on VTLN features was used. As a contrast experiment, when no SAT was used in acoustic model training, standard CMLLR was performed in recognition.

Finally, Maximum Likelihood Linear Regression (MLLR) (Leggetter and Woodland, 1995a) was applied to the means of the acoustic model in recognition. A regression class tree (Leggetter and Woodland, 1995b) was used to adjust the number of regression classes to the amount of data available.

Since both CMLLR and MLLR are text dependent, a two pass setup is needed. Also, since CMLLR is carried out in a speaker dependent manner, and since no speaker identities were provieded in the evaluation, an automatic speaker labeling was done. For SAT the speaker labels provided in the training data was used. The details of the two-pass system is described in Sec. 5.2.

3.3 Discriminative Training

To refine the ML trained acoustic model discriminative training was performed. Here the Maximum Mutual Information (MMI) criterion and the Minimum Classification Error (MCE) criterion were used as they have proven to perform best in our system. For the experiments the lattice based MCE was taken, which was originally presented in (Macherey et al., 2005) for a large vocabulary speech recognition task. As in ML training, only the manually transcribed training data was used. The discriminative training was initialized with the ML trained acoustic model. The word-conditioned word lattices used in training were generated with the VTLN/voicing system in combination with a bigram language model. For MCE the spoken word sequence needs to be contained in the lattice. To guarantee this the best alignment of the spoken word sequence was merged into the training lattices. During discriminative training, we used the exact match approach for acoustic rescoring, i.e., the word boundary times were kept fixed, and a unigram language model (Schlüter et al., 2001). The optimal number of training iterations was determined by a recognition on the development corpus and was about 10. The resulting models comprise about 800–900k Gaussians and have about 10% fewer densities than the corresponding ML trained models.

For completness we also ran experiments for Minimum Word Error (MWE). We used the time alignment based approximate accuracy in training (Povey, 2004) for efficiency reasons and because tests have shown that using the exact accuracy in training does not improve the recognition performance (Heigold et al., 2005). In contrast to other corpora (Macherey et al., 2005) I-smoothing was essential on the TC-STAR corpus.

4 Language Modeling

4.1 Baseline Language Modeling

The language model (LM) training also was done using the restricted task data. For English, the data includes the transcriptions of the acoustic training data and the FTE data. From both data sets we trained seperate case sensitive fourgram LMs. The applied smoothing was modified Kneser-Ney discounting with interpolation. The final LM was the result of a linear interpolation of the two preliminary models, where the interpolation weights were optimized on the English development set. We used the SRI Language Modeling Toolkit to build and interpolate the LMs (Stolcke, 2002). For English the optimal weights were 0.71 for VT and 0.29 for FTE. For Spanish additional restricted data from the Spanish Parliament (SPC) was used. Thus, three preliminary LMs were built. The linear interpolation weights were optimized in a grid search on the Spanish development data. The optimal weights were: VT: 0.53, FTE: 0.15, and SPC: 0.32. Table 3 gives the perplexities of the final LMs on the development and evaluation data.

4.2 Punctuation Modeling

For punctuation a sentence segmentation algorithm, also applied by the RWTH SLT system was used. For each estimated sentence break, a full stop is inserted; no further punctuation marks were produced. The segmentation approach originates from (Stolcke et al., 1998). A decision for placing a segment boundary is made based on a loglinear combination of language model and prosodic features. In contrast to existing approaches, an explicit optimization over the number of words in the segment is performed by adding a length model feature. For a more detailed presentation of this method, see the presentation of the RWTH spoken language translation system (Matusov et al., 2006).

5 Search Issues

5.1 Baseline One-Pass Recognizer

The RWTH baseline system realizes a one-pass four-gram Viterbi decoder using 6-state left-to-right HMM cross-word generalized triphone models (Ney et al., 1998) (Beulen et al., 1999). HMM states are tied pairwise such that each 6-state HMM is modeled by three separate Gaussian mixture distributions. A phonetic decision tree is used for tying the triphone models (Beulen et al., 1997). The size of the lexical pronunciation tree is further reduced via determinization on the Gaussian mixture densities level. We use two acceleration techniques: fast likelihood calculation (Kanthak et al., 2000) and language model look ahead (Ortmanns et al., 1996). The baseline system uses voicing features (cf. Sec. 3.1) and fast VTLN (cf. Sec. 3.2).

5.2 Two-Pass Speaker Adapted System

As described in Sec. 3.2, a two-pass search strategy is used to facilitate speaker adaptation. The first pass was performed using the baseline VTLN/voicing system, with the ML estimated acoustic model. Since no fine-grained segmentation of the data was provided in the evaluation, the complete recordings were used as input to the system. The recordings varied in length between a couple of minutes and half an hour. The silence information from the first recognition pass is used to segment the audio data for the second pass. The segment breaks are chosen at the longest silence regions in such a way that no segment is longer than 35s, while keeping the number of segments at a minimum. To provide a speaker labeling, a generalized likelihood ratio based segment clustering with a *Bayesian* information criterion based stopping condition was applied to the segmented recognition corpus (Chen and Gopalakrishnan, 1998). The segmented and clustered corpus was used to estimate the CMLLR and MLLR matrices needed by the adaptation. The second pass finally was performed using the best acoustic models, discriminatively trained on the CMLLR-SAT transformed features, and adapted using the estimated CMLLR and MLLR matrices.

5.3 System Combination

A common method to improve the recognition performance is to combine the output of several recognizers. System combination usually gives largest improvements, if the individual systems to be combined lead to similar performance and are complementary w.r.t. the errors they produce. A measure for the suitability for system combination is the oracle word error rate (WER) calculated on the results of the recognizers.

Ideally, the recognizers are derived from different acoustic and language models as well as different feature sets like MFCC and PLP (Perceptual Linear Prediction) features. The resulting systems show a sufficient variability in output and score distribution to make them suitable for combination.

Parallel development of complementary systems with comparable performance can be very time consuming. On the other hand, the development cycle of a state-of-the-art ASR system involves subsequent creation of suboptimal systems due to techniques like adaption and discriminative training. Since in the latter case the oracle WER still is considerably lower compared to the best single systems WER, the combination of these systems seems to be justified, see Table 8.

Two system combination techniques are well-known from literature: ROVER (Recognizer Output Voting Error Reduction) (Fiscus, 1997) and Confusion Network Combination (CNC) (Mangu et al., 2000) (Evermann and Woodland, 2000). We applied both techniques to the English corpus. In addition, a new combination technique based on a frame-wise word error measure (Wessel et al., 2001) was tested, which preserves both the word graph structure and the word boundaries. The new combination technique and an exhaustive comparison with ROVER and CNC will be presented in (Hoffmeister et al., 2006).

We used a modified version of ROVER where the confidence scores are weighted by a system dependent factor. We tried linear and exponential weights.

ROVER considers only the best hypthoses of each system. In contrast, CNC and frame-based combination take word graphs as input. A word graph contains much more information than a single hypothesis. We expected that the latter two combination techniques are able to take advantage of the additional information and outperform ROVER. For CNC and frame-based combination we did a weighted combination, where system dependent weights were trained on the development corpus.

5.4 Bayes Risk Minimization

The quality of a speech recognition system is typically assessed by its word error rate (WER). However, the standard decision rule is based on minimizing the Bayes risk using the *sentence* instead of the *word* error count as cost function. As a consequence, a rescoring pass using the Minimum Bayes Risk (MBR) criterion with a WER based cost function was applied. The experiments reported here were carried out with the algorithm proposed in (Stolcke et al., 1997) which is applied on N-best lists.

6 Experiments

The experiments described in this paper were done in the context of the second TC-STAR ASR evaluation campaign. To monitor the progress of the system development several recognition experiments were performed comparing the effectiveness of different methods applied. Due to the large number of available methods, not all possible combinations were investigated.

Since the evaluation data certainly was not available beforehand, not all contrast experiments carried out on the development data were performed on the evaluation data. For Spanish, the development was mainly carried out on the EPPS part of the development corpus.

6.1 Baseline System

As described in Sec. 5.1 the baseline system already included VTLN and voicing features. As a contrast, and for use in system combination, experiments were also performed with a plain baseline without VTLN and voicing. Table 4 summarizes the results comparing the two baseline systems, both for English and Spanish.

Table 4: *Baseline WER [%] on EPPS development data.*

	English	Spanish
Baseline	18.5	13.2
VTLN+voice	17.2	11.9

6.2 Speaker Adaptation

On top of the ML baseline system already including VTLN, four different adapted systems were used, differing w.r.t. SAT model, and MLLR usage. Table 5 show the performance of the different systems for English and Spanish on the development corpora used. While SAT gives a clear improvement in the case without MLLR, SAT with MLLR was not observed to lead to further improvements. On the other hand, when SAT is used the improvement of MLLR is somewhat inconclusive: for English the improvement is substantial, but for Spanish it is negligible. Note that for Spanish, the baseline already contains an improved language model.

Table 5: *Adaptation WER[%] on EPPS development data.*

	English	Spanish
Baseline	17.2	10.7
CMLLR	15.7	9.2
SAT	15.2	8.6
CMLLR+MLLR	14.0	8.6
SAT+MLLR	14.0	8.6

6.3 Discriminative Training

Table 6 summarizes the improvements resulting from discriminative training. Note that discriminative training in combination with MLLR did not perform consistently: for English discriminative training is beneficial whereas for Spanish the word error rate even increases. However, CMLLR-SAT combined with discriminative training and MLLR yields improvements on both corpora (see also discussion in Sec. 6.6). Furthermore, MCE slightly outperformed MMI.

Table 6: *Discriminative training performance (WER[%]) on EPPS development data.*

	English	Spanish
MLLR	14.0	8.6
MLLR+MMI	13.6	8.8
MLLR+MMI+SAT	13.3	-
MLLR+MCE+SAT	13.1	8.0

We also compared the performance of different training criteria. Note that the results are without MLLR, i.e., on the VTLN-vo+CMLLR-SAT system. The word error rates on the EPPS English development data are: 15.1% (ML), 14.5% (MMI), 14.3% (MCE), and 14.3% (MWE).

6.4 Bayes Risk Minimization

Table 7 compares the results for the best two-pass, SAT-based, discriminatively trained systems with and without using MBR. In English a marginal improvement was obtained whereas in Spanish no improvement could be observed. The observed reduced performance may be due to the relatively low error rates obtained for these tasks.

Table 7: *Performance of Bayes risk minimization (WER[%]) on EPPS development data.*

	English	Spanish
No MBR	12.9	7.8
MBR	12.8	7.8

6.5 System Combination

System combination results were only produced for English. We decided to use the following individual systems for system combination: CMLLR+MLLR,

CMLLR+MMI+MLLR, CMLLR-SAT+MMI+MLLR, CMLLR-SAT+MMI+MLLR+NEW-LM. The difference between the last two system is due to an optimization of the lexicon and the language model.

Table 8 summarizes the results: none of the techniques were able to yield a significant decrease in WER on the evaluation set. That was quite surprising since on the development set we got promising results and the oracle WER shows some potential for system combination on the evaluation set. For Spanish first experiments did not even show

Table 8: *Performance of system combination, WER[%].*

combiantion method	systems	WER dev	WER eval
single systems	+CMLLR/MLLR	14.1	11.8
	+MMI	13.7	11.7
	+SAT	13.3	10.8
	+new lexicon and LM	12.9	10.3
Oracle		10.8	8.6
ROVER	w/o +CMLLR/MLLR	13.0	10.5
	+ conf. scores	12.6	10.5
	+ linear weighted conf. scores	12.5	10.4
	+ exp. weighted conf. scores	12.6	10.3
CNC		13.1	10.6
	+ weights	12.9	10.2
Frame Based		12.8	10.7
	+ weights	12.5	10.3

a decrease in WER on the development set, so we decided not to try internal system combination for the evaluation.

6.6 Summary of Results

Tables 9 and 10 show the chronological progression of the results during the preparation for the evaluation campaign[1], as well as the corresponding results for the evaluation corpus, where available. Note that while the separate improvements of STM-SAT and discriminative training were small, the combined improvement was larger than the sum of the separate improvements, when compared to a ML trained system with both CMLLR and MLLR. A similar effect has been described in (Povey, 2004), where discriminative training was reported to give larger improvements when the system is using SAT and MLLR, as compared to only using MLLR.

Table 9: *Overview of English system performance (WER[%]).*

	Dev	Eval
Baseline	18.5	-
+VTLN+voice	17.2	14.4
+CMLLR	15.7	-
+MLLR	14.0	11.8
+MMI	13.6	11.7
+SAT	13.3	10.8
+New LM	12.9	10.3
+MBR	12.8	10.2

[1]The entry *Tuning* refers to language model scale tuning

Table 10: *Overview of Spanish system performance (WER[%]). Note that the evaluation data contains EPPS and STC data.*

	Dev	Eval
Baseline	13.2	-
+VTLN+voice	11.9	-
+New LM	10.7	16.1
+MLLR	8.6	11.3
+MCE	8.8	11.1
+SAT	8.0	-
+Tuning	7.8	10.2

7 Conclusions & Outlook

In this work, the RWTH automatic speech recognition systems developed for the second TC-STAR evaluation campaign 2006 were presented. The systems were designed to transcribe parliamentary speeches taken from the European Parliament Plenary Sessions (EPPS) in European English and Spanish, as well as speeches from the Spanish Parliament. Using a two-pass decoding strategy a number of improvements could be obtained. Using several speaker adaptation and normalization schemes, speaker adaptive training, MCE and MMI training, and *Bayes* risk minimization, the overall improvement obtained on top of the baseline system ranged between 30% and 40% relative WER reduction. For all relevant system components, contrastive results are presented on the EPPS Spanish and English data. In addition experiments on system combination were performed but not used in the final evaluation.

Acknowledgements

This work was funded by the European Union under the Human Language Technologies project TC-STAR (FP6-506738).

8 References

K. Beulen, E. Bransch, and H. Ney. 1997. State-tying for context dependent phoneme models. In *Proc. European Conf. on Speech Communication and Technology*, pages 1179 – 1182, Rhodes, Greece.

K. Beulen, S. Ortmanns, and C. Elting. 1999. Dynamic programming search techniques for across-word modelling in speech recognition. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, pages 609 – 612, Phoenix, AZ, USA.

M. Bisani and H. Ney. 2003. Multigrambased grapheme-to-phoneme conversion for LVCSR. In *Proc. European Conf. on Speech Communication and Technology*, volume 2, pages 933 – 936, Geneva, Switzerland.

M. Bisani and H. Ney. 2005. Open vocabulary speech recognition with flat hybrid models. In *Proc. European Conf. on Speech Communication and Technology*, pages 725 – 728, Lisbon, Portugal.

S. S. Chen and P. S. Gopalakrishnan. 1998. Clustering via the Bayesian information criterion with applications in speech recognition. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, volume 2, pages 645 – 648.

E. Eide and H. Gish. 1996. A parametric approach to vocal tract length normalization. In *Proc. IEEE Int. Conf.*

on Acoustics, Speech, and Signal Processing, volume 1, pages 346 – 349, Atlanta, GA.

http://www.europarl.eu.int. European Parliament.

G. Evermann and P.C. Woodland. 2000. Posterior probability decoding, confidence estimation and system combination. In *In Proc. Speech Transcription Workshop*, College Park, MD, USA.

J. G. Fiscus. 1997. A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (ROVER). In *Proc. IEEE Automatic Speech Recognition and Understanding Workshop*.

D. Giuliani G. Stemmer, F. Brugnara. 2005. Adaptive training using simple target models. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, volume 1, pages 997 – 1000, Philadelphia, PA, USA.

M. J. F. Gales. 1998. Maximum likelihood linear transformations for HMM-based speech recognition. *Computer Speech and Language*, 12(2):75 – 98.

G. Heigold, W. Macherey, R. Schlüter, and H. Ney. 2005. Minimum exact word error training. In *Proc. IEEE Automatic Speech Recognition and Understanding Workshop*, pages 186 – 190, San Juan, Puerto Rico.

B. Hoffmeister, T. Klein, R. Schlüter, and H. Ney. 2006. Frame based system combination and a comparison with weighted ROVER and CNC. In *Proc. Int. Conf. on Spoken Language Processing*, Submitted.

S. Kanthak, K. Schütz, and H. Ney. 2000. Using SIMD instructions for fast likelihood calculation in LVCSR. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, pages 1531 – 1534, Istanbul, Turkey.

http://www.lc-star.com. LC-STAR, Lexica and Corpora for Speech-to-Speech Translation Components.

C. J. Leggetter and P. C. Woodland. 1995a. Maximum likelihood linear regression for speaker adaptation of continuous density hidden markov models. *Computer Speech and Language*, 9(2):171 – 185.

C.J. Leggetter and P.C. Woodland. 1995b. Flexible speaker adaptation using maximum likelihood linear regression. In *Proc. ARPA Spoken Language Technology Workshop*, pages 104 – 109, Austin, TX, USA.

W. Macherey, L. Haferkamp, R. Schlüter, and H. Ney. 2005. Investigations on error minimizing training criteria for discriminative training in automatic speech recognition. In *Proc. European Conf. on Speech Communication and Technology*, pages 2133 – 2136, Lisbon, Portugal.

L. Mangu, E. Brill, and A. Stolcke. 2000. Finding consensus in speech recognition: Word error minimization and other applications of confusion networks. *Computer Speech and Language*, 14(4):373 – 400.

E. Matusov, R. Zens, D. Vilar, A. Mauser, M. Popović, S. Hasan, and H. Ney. 2006. The RWTH machine translation system. In *Proc.* TC-STAR *Evaluation Workshop*, In preparation.

H. Ney, L. Welling, S. Ortmanns, K. Beulen, and F. Wessel. 1998. The rwth large vocabulary continuous speech recognition system and spoken document retrieval. In *Proc. 24th Annual Conference of the IEEE Industrial Electronics Society*, pages 2022 – 2027, Aachen, Germany.

S. Ortmanns, H. Ney, and A. Eiden. 1996. Language-model look-ahead for large vocabulary speech recognition. In *Proc. Int. Conf. on Spoken Language Processing*, pages 1095 – 2098, Philadelphia, PA, USA.

D. Povey. 2004. *Discriminative Training for Large Vocabulary Speech Recognition*. Ph.D. thesis, Cambridge, England.

R. Schlüter, W. Macherey, B. Müller, and H. Ney. 2001. Comparison of discriminative training criteria and optimization methods for speech recognition. *Speech Communication*, 34:287 – 310.

A. Stolcke, Y. König, and M. Weintraub. 1997. Explicit word error minimization in N-best list rescoring. In *Proc. European Conf. on Speech Communication and Technology*, pages 373 – 400, Rhodes, Greece.

A. Stolcke, E. Shriberg, R. Bates, M. Ostendorf, D. Hakkani, M. Plauche, G. Tür, and Y. Lu. 1998. Automatic detection of sentence boundaries and disfluencies based on recognized words. In *Proc. Int. Conf. on Spoken Language Processing*, Sidney, Australia.

A. Stolcke. 2002. SRILM - An extensible language modeling toolkit. In *Proc. Int. Conf. on Spoken Language Processing*, Denver, CO, USA.

http://www.tc-star.com. TC-STAR, Technology and Corpora for Speech-to-Speech Translation Components.

L. Welling, S. Kanthak, and H. Ney. 1999. Improved methods for vocal tract normalization. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, pages 761 – 764, Phoenix, AZ, USA.

F. Wessel, R. Schluter, and H. Ney. 2001. Explicit word error minimization using word hypothesis posterior probabilities. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, volume 1, pages 33 – 36, Salt Lake City, Utah, USA.

A. Zolnay, R. Schlüter, and H. Ney. 2002. Robust speech recognition using a voiced-unvoiced feature. In *Proc. Int. Conf. on Spoken Language Processing*, volume 2, pages 1065 – 1068, Denver, CO, USA.