# ENGLISH BROADCAST NEWS SPEECH RECOGNITION BY HUMANS AND MACHINES

Samuel Thomas, Masayuki Suzuki\*, Yinghui Huang, Gakuto Kurata\*, Zoltan Tuske, George Saon, Brian Kingsbury, Michael Picheny

\*IBM Research AI, Tokyo, Japan IBM Research AI, Yorktown Heights, USA

#### ABSTRACT

With recent advances in deep learning, considerable attention has been given to achieving automatic speech recognition performance close to human performance on tasks like conversational telephone speech (CTS) recognition. In this paper we evaluate the usefulness of these proposed techniques on broadcast news (BN), a similar challenging task. We also perform a set of recognition measurements to understand how close the achieved automatic speech recognition results are to human performance on this task. On two publicly available BN test sets, DEV04F and RT04, our speech recognition system using LSTM and residual network based acoustic models with a combination of n-gram and neural network language models performs at 6.5% and 5.9% word error rate. By achieving new performance milestones on these test sets, our experiments show that techniques developed on other related tasks, like CTS, can be transferred to achieve similar performance. In contrast, the best measured human recognition performance on these test sets is much lower, at 3.6% and 2.8% respectively, indicating that there is still room for new techniques and improvements in this space, to reach human performance levels.

*Index Terms*— Broadcast News, Automatic Speech Recognition, Deep neural networks.

## 1. INTRODUCTION

Prior to the recent ubiquitous deployment of automatic speech recognition technology for various device user interfaces, two key domains of interest for application of automatic speech recognition technology were conversational telephone speech (CTS) and broadcast news (BN). Interest in these domains was primarily fueled by various DARPA programs [1]. More recently, by employing various deep learning techniques, performance of speech recognition systems on the CTS task is getting close to human parity. Several sites have made significant progress to lower the WER to within the 5%-10% range on the Switchboard-CallHome subsets of the Hub5 2000 evaluation [2, 3, 4, 5]. Given the progress on conversational telephone speech, we focus on the other closely related broadcast news recognition task that received similar attention within the DARPA EARS program. One of the key objectives of this study is to understand how deep learning based techniques developed on CTS generalize to the BN task.

In the BN domain, speech recognition systems need to deal with wide-band signals collected over a wide variety of speakers with different speaking styles, in various background noise conditions, and speaking on a wide variety of news topics. Most of the speech is well articulated and is formed similarly to written English. In contrast, CTS is spontaneous speech recorded over a telephone channel Tom Dibert, Alice Kaiser-Schatzlein, Bern Samko

Appen, Sydney, Australia



Fig. 1. The NIST STT Benchmark Test History - May'09 [7]

that introduces additional artifacts in addition to numerous speaking styles. Conversational speech is interspersed with portions of overlapping speech, interruptions, restarts and back-channel confirmations between participants. In terms of the amount of training data available from the DARPA EARS program for training systems on CTS and BN, there are a few significant differences as well. The CTS acoustic training corpus consists of approximately 2000 hours of speech with human transcriptions [2]. On the other hand, for the BN task, only about 140 hours of data is carefully transcribed. The remaining ~9000 hours of available speech are TV shows with closed captions. In other words, models being developed for BN typically use lightly supervised transcripts for training [6].

The EARS program led to significant advances in speech recognition technology for both domains, with the development of techniques that could ingest large quantities of unsupervised or semisupervised training data, discriminative training of recognition models, methods to deal with channel and speaker variabilities in the data, real-time decoding of test data, and also approaches to combine outputs from various systems [8, 9, 10, 11]. Several of these techniques have further been extended to build ASR systems on broadcast news data in various languages [12, 13, 14]. Figure 1 shows progress made in this domain over the past two decades. More recently, as part of the MGB Challenge, in addition to the core ASR problem, several other related tasks - speaker diarization and lightly supervised alignment of data have also been studied [15, 16].

In [2, 3] we describe state-of-the-art speech recognition systems on the CTS task using multiple LSTM and ResNet acoustic models trained on various acoustic features along with word and character LSTMs and convolutional WaveNet-style language models. This advanced recipe achieves 5.1% and 9.9% on the Switchboard and Call-Home subsets of the Hub5 2000 evaluation. In this paper we develop a similar but simpler variant for BN. As described earlier, by developing this system we investigate how these earlier proposed systems can be trained on BN data which are not human annotated but are created using closed captions. To create these systems, instead of adding all the available training data we carefully select a reliable subset. We then train LSTM and residual network based acoustic models with a combination of n-gram and neural network language models on this selected data. In addition to automatic speech recognition results, similar to [2], we also present human performance on the same BN test sets. These evaluations allow us to properly benchmark our automatic system performance. Similar to earlier human performance evaluations on CTS, we observe a significant gap between human and automatic results.

The rest of the paper is organized as follows. In Section 2 we describe the human evaluation experiments on two broadcast news test sets - RT04 and DEV04F. We also compare the recognition errors we observe with human and automatic recognition systems. Section 3 describes the development of our ASR systems - training data selection, acoustic and language model building. In Section 4 we present WER results using the proposed system. The paper concludes with a discussion in Section 5.

## 2. HUMAN TRANSCRIPTION EXPERIMENTS

Similar to [2], human performance measurements on two broadcast news tasks - RT04 and DEV04F - are carried out by Appen. For these evaluations we limit the audio from the test sets to only regions of speech that are marked for scoring using the original references and scoring scripts provided during the EARS evaluation. After processing, the RT04 test set has 4 hours of BN data from 12 shows with about 230 overlapping speakers across the shows. The DEV04F test set is smaller, with about 2 hours of data from 6 shows with close to 100 overlapping speakers across the shows.

The first round of transcripts was produced by three independent transcribers, followed by quality checking by a fourth senior transcriber. All four transcribers are native US English speakers and were selected based on the quality of their work on past transcription projects. The transcriptions were produced in line with LDC transcription guidelines for hyphenations, spelled abbreviations, contractions, partial words, non-speech sounds, etc. that were used to produce the original transcripts for these test sets. The three primary transcribers took 14-16 times real-time (xRT) for the first pass followed by an additional 3xRT for the second quality checking pass (by Transcriber 4). Both passes involved listening to the audio multiple times: around 3-4 times for the first pass and 1-2 times for the second. In order to use NIST's scoring tool, sclite [17], the human annotations were converted into CTM files which have time-marked word boundary information. The transcriptions were also filtered to remove non-speech markers, partial words, punctuation marks etc as described in [2]. Table 1 shows the error rates of the three transcribers after quality checking by the fourth transcriber.

	DEV04F	RT04
Transcriber 1	4.4	3.6
Transcriber 2	4.4	3.2
Transcriber 3	3.6	2.8

Table 1. Human Performance (WER%) on RT04 and DEV04F.

Compared to the human transcription results on the CTS tasks,

5.1% and 6.8% on the Switchboard and CallHome subsets of the Hub5 2000 evaluation [2], the word error rate on BN is much lower. Although this reduction could be because BN speech is well articulated, the transcribers reported that these test sets were much denser with respect to speech content, had considerable background noise, and a significant number of named entities that required lookup to ensure correctness, compared to traditional CTS test sets. The best WER results we obtain, 3.6% and 2.8%, also fit in the expected human transcription error range indicated in Figure 1. A more detailed error analysis and comparison of human and automatic recognition is presented in the Discussion section.

## 3. ASR SYSTEM BUILDING

As described earlier, one differentiating characteristic of ASR system builds for this BN task is the limited amount of carefully annotated manual transcriptions. Prior to the EARS program, LDC released about 144 hours of careful manual annotations for a portion of the Hub4 acoustic training data collected between May 1996 and January 1998. In addition to this, several sources of BN data were available for training acoustic models during the EARS program period with just closed caption transcripts. These data sources include about 1000 hours of data as part of different data releases collected between 1998-2001 (TDT2 and TDT4) and about 7000 hours of broadcast news released in 2003 as part of the EARS program (BN03). In this paper we use processed versions of these data sources to build deep neural network based acoustic and language models.

## 3.1. Training Data Preparation

To process the BN data with noisy closed captions, the data is first decoded with multiple off-the-shelf broadband ASR systems using a biased LM created with the available closed captions. Based on the confidence scores of these decodes and agreement between the multiple system decodes, we perform a strict filtering of the data to create 2 sub-corpora that we consider have very reliable transcripts -

- The BN-400 Corpus This is a corpus of about 430 hours of broadcast news data selected from the data sources described above. This data corpus includes 144 hours of carefully transcribed audio along with data with semi-supervised transcripts created via a biased decode of the matching audio.
- The BN-1300 Corpus This corpus is an extended version of the BN-400 corpus with about 900 additional hours of broadcast news.

#### 3.2. Acoustic Model Development

As discussed earlier, one of the key objectives of this work is to verify the usefulness of our earlier proposed system strategy for CTS. In [2], two kinds of acoustic models, a convolutional and a non-convolutional acoustic model with comparable performance, are used since they produce good complementary outputs which can be further combined for improved performance. The convolutional network used in that work is a residual network (ResNet) and an LSTM is used as the non-convolutional network. The acoustic scores of these systems are subsequently combined for the final decodes. Similar to that work, in this paper also we train ResNet and LSTM based acoustic models.

Both these acoustic models are based on speaker transformed features. The ResNet uses 40 dimensional VTL-warped log-mel fea-

Input	$3 \times 40 \times 76$
	conv 5x5, 64;
Stage 0	maxpool $(2 \times 1)$
	initStride $1 \times 1$ ;
Stage 1	$3 \times [\text{conv } 3 \times 3, 64 \text{ feat. maps, conv } 3 \times 3, 64 \text{ feat. maps}]$
	initStride $2 \times 1$ ;
Stage 2	$3 \times [\text{conv } 3 \times 3, 128 \text{ feat. maps, conv } 3 \times 3, 128 \text{ feat. maps}]$
	initStride $2 \times 1$ ;
Stage 3	$3 \times [\text{conv } 3 \times 3, 256 \text{ feat. maps, conv } 3 \times 3, 256 \text{ feat. maps}]$
	initStride $2 \times 2$ ;
	$3 \times [\text{conv } 3 \times 3, 512 \text{ feat. maps, conv } 3 \times 3, 512 \text{ feat. maps}];$
Stage 4	maxpool $(2 \times 2)$
Output	3 × FC 2084; FC 1024; FC 32K

Table 2. ResNet acoustic model architecture

tures along with their  $\Delta$  and  $\Delta\Delta$  transforms. The LSTM based network is trained on 40 dimensional FMLLR features appended with 100 dimensional ivectors and 40 dimensional VTL-warped log-mels along with their  $\Delta$  and  $\Delta\Delta$  parameters. The speaker transformed features - FMLLR and VTL-warped features and ivectors, are derived using traditional HMM-GMM based systems trained on the BN-400 corpus. We model the acoustic space with 32K context-dependent HMM states.

Similar to the architecture in [2], we train an LSTM acoustic model with 6 bidirectional layers having 1024 cells per layer (512 per direction), one linear bottleneck layer with 256 units and an output layer with 32K units corresponding to the contextdependent HMM states we derived in the HMM-GMM system build. The model is trained using non-overlapping subsequences of 21 frames. Subsequences from different utterances are grouped into mini-batches of size 128 for processing speed and reliable gradient estimates. After the cross-entropy based training on the BN-1300 Corpus has converged we also sequence train the model using the 144 hours of carefully transcribed audio.

To complement the LSTM acoustic model, we train a deep Residual Network based on the best performing architecture proposed in [2]. The ResNet has 12 residual blocks followed by 5 fully connected layers. To effectively train this network with 25 convolutional layers, a short-cut connection is placed between each residual block to allow for additional flow of information and gradients. Each layer has a batch normalization layer as well. Table 2 gives a summary of the network architecture of the ResNet model. The ResNet consists of several stages with different numbers of feature maps in each stage: 64 in stage 1, 128 in stage 2, 256 in stage 3 and 512 in stage 4. Each stage has an "initStride" which indicates the (frequency  $\times$  time) stride for the first block of that stage as the number of feature maps is increased. The stride applies to both the first  $3 \times 3$  convolution of each block and also the  $1 \times 1$  convolution in projection shortcut between each block. The ResNet acoustic model is trained using the cross-entropy training criterion on the BN-1300 Corpus and then sequence trained using the 144 hours of carefully transcribed audio.

#### 3.3. Language Model Development

Similar to the development of acoustic models, several kinds of ngram and neural network based language models are built on this BN task. For the initial decode that produces word lattices, an n-gram and a feed forward neural network language model are first built. To rescore the word lattices and n-best lists produced by these models, advanced LSTM based NN language models are also constructed.

AM	LM	DEV04F	RT04
LSTM	n-gram	7.6	7.7
ResNet	n-gram	9.6	8.9
LSTM	n-gram + FFNN-LM	7.2	7.0
ResNet	n-gram + FFNN-LM	9.0	8.1
LSTM+ResNet	n-gram + FFNN-LM	7.2	7.0

Table 3. ASR decoding results (WER%) on RT04 and DEV04F.

	DEV04F	RT04
LSTM1-LM rescoring	6.6	6.1
LSTM2-LM rescoring	6.6	6.1
LSTM1/LSTM2-LM rescoring	6.5	5.9

Table 4. LSTM rescoring results (WER%) on RT04 and DEV04F.

The primary language model training text for all these models consists of a total of 350M words from different publicly available sources released by LDC during the GALE [1] and EARS evaluation periods suitable for broadcast news. The baseline language model is a linear interpolation of word 6-gram models, one for each corpus with a vocabulary size of about 80K words. We train a feed forward neural network model based on the same data and vocabulary as the n-gram language model described above. The neural network model (FFNN-LM) uses an embedding size of 120, a hidden layer size of 1200 and the maxout non-linearity [18]. We use noise contrastive estimation to train this unnormalized NNLM [19]. For decoding experiments, the FF-NNLM is interpolated with the baseline 6-gram arpabo with an interpolation weight set to 0.5.

In addition to the n-gram and feed forward neural network language models, we also train two different flavors of LSTM language models with the same vocabulary and training data as described above. The first LSTM model (LSTM1-LM) consists of one word embedding layer with 256 units, four LSTM layers with 1024 units, one fully-connected layer, and one softmax layer. The second to fourth LSTM layers and the fully-connected layer allow residual connections [20]. Dropout is applied to the vertical dimension only and not applied to the time dimension. We trained this model to minimize the cross-entropy objective using Adam for learning rate control [21]. The second LSTM based LM (LSTM2-LM) consists of two LSTM layers, each layer with 2048 nodes and a word embedding size of 512. Before the softmax-based estimation of an 80Kdimensional posterior vector, the feature space was reduced to 128 by a linear bottleneck layer. During the training various dropout techniques were applied [22]. First, the outputs of the embedding and each LSTM layer were masked at a 10% rate. Second, 10% dropout was also applied on the embedding weights, and also on the parameters of the recurrent connection of the LSTMs. These weight masks were kept constant during processing a mini-batch of sequences. In the final step of the training, the model was fine-tuned on the best matching resource, the EARS BN data. The SGD based model training uses a batch size of 128 and a Nesterov momentum of 0.9 to optimize model parameters on the cross-entropy criterion.

#### 4. ASR EXPERIMENTS AND RESULTS

The acoustic and language models described above are used to decode the RT04 and DEV04F test sets. We use the same speech segments that were provided to the human annotators for our various experiments. In our first set of experiments we separately test the

	DEV04F		RT04		
	ASR	Human	ASR	Human	
Sub	3.2	1.9	3.1	1.6	
Del	2.2	0.8	2.2	0.6	
Ins	1.1	0.9	0.6	0.6	
All	6.5	3.6	5.9	2.8	

 Table 5.
 Overall substitution, deletion and insertion errors of humans and ASR system.

LSTM and ResNet models in conjunction with the n-gram and FF-NNLM, before combining scores from the two acoustic models. Table 3 shows the individual and combined results we obtain on both the test sets. In comparison with the results obtained on the CTS evaluation with similar acoustic models [2], the LSTM and ResNet operate at similar WERs. Unlike results observed on the CTS task, no significant reduction in WER is obtained after scores from both the LSTM and ResNet models are combined. The LSTM model with an n-gram LM individually performs quite well and its results further improve with the addition of the FF-NNLM.

For our second set of experiments word lattices are generated after decoding with the LSTM+ResNet+n-gram+FF-NNLM model. We generated n-best lists from these lattices and rescored them with the LSTM1-LM. LSTM2-LM is also used to rescore word lattices independently. Table 4 shows the results after our rescoring experiments. We observe significant WER gains after using the LSTM LMs similar to those reported in [2]. By rescoring outputs with both LSTM1 and LSTM2, we achieve new performance milestones with final WERs of 6.5% and 5.9% on DEV04F and RT04 respectively.

Our ASR results have clearly improved state-of-the-art performance on these test sets compared to the various results reported in [8, 9, 23]. Significant progress has also been made compared to systems developed over the last decade, as shown in Figure 1.

#### 5. DISCUSSIONS

When compared to the human performance results, the absolute ASR WER is about 3% worse. From Table 5 we observe that although the machine and human insertion error rates are comparable, the ASR system has much higher substitution and deletion error rates. Tables 6 and 7 list the 10 most frequent errors of each type. We draw the following observations based on these errors -

- 1. There is a significant overlap in the words that ASR and humans delete, substitute and insert.
- 2. Humans seem to be careful about marking hesitations %hesitation is the most inserted symbol. Hesitations seem to be important in conveying meaning to the sentences in human transcriptions. The ASR systems however focus on blind recognition and not in improving the meaning with appropriate pauses, etc. To measure the extent of this process, we score the human transcripts without any hesitations in a separate experiment and observe a 0.1% absolute improvement in WER.
- 3. Machines have trouble recognizing short function words {*the, and, of, a, that*} and these get deleted the most. Humans, on the other hand, seem to catch most of them. It is likely that these words are not fully articulated and hence the machine fails to recognize them while humans are able to infer these words even without full acoustic evidence since they may have a better model of syntax/semantics.

DEV	04F	RT04		
ASR Human		ASR	Human	
8: and / in	5: too / to	21: the / a	19: the / a	
7: had / have	4: is / has	16: and / in	15: and / in	
3: a / the	4: a / the	15: a / the	11: in / and	
3: has / is	4: had / have	14: has / is	7: (%hes) / a	
3: on / in	3: the / (%hes)	11: (%hes) / a	6: this / the	
3: that / it	3: on / in	11: in / and	5: are / were	
3: too / to	3: (%hes) / a	8: that / it	4: and / then	
3: this / the	3: in / and	7: this / the	4: as / is	
2: (%hes) / and	3: and / in	6: is / as	3: (%hes) / and	
2: and / an	2: are / were	4: it / that	2: an / in	

 Table 6. Most frequent substitution errors for humans and ASR systems on DEV04F and RT04

Deletions			Insertions				
DEV	DEV04F RT04		DEV04F		RT04		
ASR	Human	ASR	Human	ASR	Human	ASR	Human
49: the	13: the	92: and	21: and	15: and	33: (%hes)	10: i	31: (%hes)
43: and	10: a	69: the	21: the	12: colon	14: and	8: and	13: a
21: a	8: and	52: a	17: a	10: to	10: the	7: it	11: and
17: that	8: that	47: it	14: is	8: call	7: it	6: that	10: c
17: to	6: it	35: that	9: in	8: the	5: is	6: to	9: have
16: it	5: i	33: is	9: that	7: a	4: a	4: a	8: the
16: of	5: of	32: in	8: are	7: ask	4: that	4: post	6: it
15: you	5: you	28: you	8: i	5: are	4: to	3: are	5: post
14: are	4: are	22: of	8: of	4: how	3: are	3: had	5: to
10: in	3: have	19: i	7: (%hes)	3: be	3: c	3: he	4: are

**Table 7.** Most frequent deletion and insertion errors for humans andASR systems on DEV04F and RT04

- Compared to the telephone conversation confusions recorded in [2] - one symbol that is clearly missing is the back-channel response - this is probably from the very nature of the BN domain.
- Similar to telephone conversation confusions reported in [2], humans performance is much higher because the number of deletions is significantly lower - compare 2.3% vs 0.8%/0.6% for deletion errors in Table 5.

## 6. CONCLUSION

We have presented recent improvements on broadcast news transcription based on earlier established techniques shown to be useful on CTS. Our experiments on BN show that these techniques can be transferred across domains to provide highly accurate transcriptions. For both acoustic and language modeling we have demonstrated the effectiveness of LSTM and ResNet based models. To verify the extent of the improvements obtained, human evaluation experiments are also performed on the two test sets of interest. We show that there still exists a significant gap between human and machine performance and demonstrate the need for continued research on broadcast news.

#### 7. REFERENCES

- Joseph Olive, Caitlin Christianson, and John McCary, Handbook of natural language processing and machine translation: DARPA Global Autonomous Language Exploitation, Springer Science & Business Media, 2011.
- [2] George Saon, Gakuto Kurata, Tom Sercu, Kartik Audhkhasi, Samuel Thomas, Dimitrios Dimitriadis, Xiaodong Cui, Bhuvana Ramabhadran, Michael Picheny, Lynn-Li Lim, et al., "English conversational telephone speech recognition by humans and machines," arXiv preprint arXiv:1703.02136, 2017.
- [3] Gakuto Kurata, Bhuvana Ramabhadran, George Saon, and Abhinav Sethy, "Language modeling with highway LSTM," in Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE. IEEE, 2017, pp. 244–251.
- [4] Wayne Xiong, Lingfeng Wu, Fil Alleva, Jasha Droppo, Xuedong Huang, and Andreas Stolcke, "The Microsoft 2017 conversational speech recognition system," in 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2018, pp. 5934–5938.
- [5] Kyu J Han, Akshay Chandrashekaran, Jungsuk Kim, and Ian Lane, "The CAPIO 2017 conversational speech recognition system," arXiv preprint arXiv:1801.00059, 2017.
- [6] Lori Lamel, Jean-Luc Gauvain, and Gilles Adda, "Lightly supervised acoustic model training," in ASR2000-Automatic Speech Recognition: Challenges for the new Millenium ISCA Tutorial and Research Workshop (ITRW), 2000.
- [7] NIST, "NIST Rich Transcription Evaluation," https://www.nist.gov/itl/iad/mig/ rich-transcription-evaluation.
- [8] Mark JF Gales, PC Woodland, Ho Yin Chan, David Mrva, Rohit Sinha, Sue E Tranter, et al., "Progress in the CU-HTK broadcast news transcription system," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1513–1525, 2006.
- [9] Spyridon Matsoukas, J-L Gauvain, Gilles Adda, Thomas Colthurst, Chia-Lin Kao, Owen Kimball, Lori Lamel, Fabrice Lefevre, Jeff Z Ma, John Makhoul, et al., "Advances in transcription of broadcast news and conversational telephone speech within the combined EARS BBN/LIMSI system," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1541–1556, 2006.
- [10] Stanley F Chen, Brian Kingsbury, Lidia Mangu, Daniel Povey, George Saon, Hagen Soltau, and Geoffrey Zweig, "Advances in speech transcription at IBM under the DARPA EARS program," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1596–1608, 2006.
- [11] Andreas Stolcke, Barry Chen, Horacio Franco, Venkata Ramana Rao Gadde, Martin Graciarena, Mei-Yuh Hwang, Katrin Kirchhoff, Arindam Mandal, Nelson Morgan, Xin Lei, et al., "Recent innovations in speech-to-text transcription at SRI-ICSI-UW," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1729–1744, 2006.
- [12] Rohit Sinha, Mark JF Gales, DY Kim, X Andrew Liu, Khe Chai Sim, and Philip C Woodland, "The CU-HTK Mandarin broadcast news transcription system," in Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on. IEEE, 2006, vol. 1, pp. I–I.

- [13] Xin Lei, Manhung Siu, Mei-Yuh Hwang, Mari Ostendorf, and Tan Lee, "Improved tone modeling for Mandarin broadcast news speech recognition," in *Ninth International Conference* on Spoken Language Processing, 2006.
- [14] Bing Xiang, Kham Nguyen, Long Nguyen, Richard Schwartz, and John Makhoul, "Morphological decomposition for Arabic broadcast news transcription," in Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on. IEEE, 2006, vol. 1, pp. I–I.
- [15] Peter Bell, Mark JF Gales, Thomas Hain, Jonathan Kilgour, Pierre Lanchantin, Xunying Liu, Andrew McParland, Steve Renals, Oscar Saz, Mirjam Wester, et al., "The MGB challenge: Evaluating multi-genre broadcast media recognition," in Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on. IEEE, 2015, pp. 687–693.
- [16] Ahmed Ali, Stephan Vogel, and Steve Renals, "Speech recognition challenge in the wild: Arabic MGB-3," in Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE. IEEE, 2017, pp. 316–322.
- [17] NIST, "NIST Tools," https://www.nist.gov/itl/ iad/mig/tools.
- [18] Yinghui Huang, Abhinav Sethy, and Bhuvana Ramabhadran, "Fast neural network language model lookups at n-gram speeds," in *Interspeech 2017*. ISCA, 2017.
- [19] Abhinav Sethy, Stanley Chen, Ebru Arisoy, and Bhuvana Ramabhadran, "Unnormalized exponential and neural network language models," in *ICASSP*, 2015.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [21] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [22] Stephen Merity, Nitish Shirish Keskar, and Richard Socher, "Regularizing and optimizing LSTM language models," in *International Conference on Learning Representations (ICLR)*, 2018.
- [23] Tara N Sainath, Brian Kingsbury, Abdel-rahman Mohamed, George E Dahl, George Saon, Hagen Soltau, Tomas Beran, Aleksandr Y Aravkin, and Bhuvana Ramabhadran, "Improvements to deep convolutional neural networks for LVCSR," in Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on. IEEE, 2013, pp. 315–320.