Short Term Context-based Fast Multilingual Acoustic Model for Low Resource Languages

Bipasha Sen[1*], *Aditya Agarwal*[1*], *Mirishkar Sai Ganesh*[2], *Anil Kumar Vuppala*[2]

[1]Microsoft India
[2]International Institute of Information Technology, Hyderabad

bise@microsoft.com, agarwal.aditya@microsoft.com, mirishkar.ganesh@research.iiit.ac.in,
anil.vuppala@iiit.ac.in

Abstract

Multilingual automatic speech recognition (ASR) systems have led to a major step forward towards building robust ASR systems for languages with low resource availability by increasing coverage for individual languages. State of the art multilingual systems are developed with sequential networks such as recurrent neural networks (RNNs) to capture long term temporal dependencies. Training and inference in such sequential models are computationally expensive, which poses a significant challenge in terms of scalability and real-time applications. In this paper, an alternate architecture based on short term contextual temporal features learned on convolutional neural networks (CNNs) with a non-sequential discriminative network is proposed. Three low resource Indic languages, Gujarati, Tamil, and Telugu are used to ascertain that our proposed architecture trains $5.5\times$ faster and reduces the inference time by a factor of 26 while maintaining comparable word error rates (WERs) against our baseline RNN.

Index Terms: Multilingual ASR, Low Resource, CNN-DNN

1. Introduction

Multilingual automatic speech recognition (ASR) system is a single entity capable of transcribing speech utterances for multiple languages with a shared phone space. Multilingual ASR systems have seen tremendous growth and have proven to be a viable solution for building robust speech recognition systems often outperforming the monolingual counterparts, especially for languages with low resource [1–5].

Conventionally, an ASR system is a modular system comprising of multiple sequential components. Each of these components is optimized individually. The first component extracts state of the art handcrafted features such as Mel frequency cepstral coefficients (MFCCs), perceptual linear prediction (PLP) from raw speech signals. Additionally, forced alignment techniques are used to generate output phone labels for these extracted features. The second component, an acoustic model trained on the extracted feature - alignment label pairs, estimates the phone probabilities for given acoustic representations. The third component, word reconstruction, takes in the estimated phone probabilities as input and predicts a sequence of words using a language model.

Recently, there have been substantial developments on end to end (E2E) acoustic modeling which combines multiple components to optimize the system as a single unit. Techniques based on attention networks [6, 7], connectionist temporal classification (CTC) [8–10] transcribe acoustic frames into phones without the need of any predefined alignments. Techniques

---

*Bipasha and Aditya contributed equally to this work.

based on convolutional neural networks (CNNs) transcribe raw speech utterances directly into phones, eliminating the step of hand-crafted feature extraction [11–16].

In recent years, the use of speech recognition systems has soared across the globe. However, such systems are still largely limited to only a handful of languages [17]. It thus becomes imperative to build robust and scalable techniques that can be used to quickly bootstrap ASR systems for multiple languages.

The discriminative networks in the acoustic models are largely built on sequential networks such as recurrent neural networks (RNNs) [18–20]. Sequential networks have the inherent property of capturing long term contextual dependencies in a sequence. Fields like natural language processing, machine translation are largely dependent on such sequential architectures. Such networks are however computationally expensive due to limited scope in parallelization. Deep neural network (DNN) on the other hand is a non-sequential architecture, providing computational efficiency but at a cost of high WERs [12].

Phonemes have very short contextual dependencies. For instance, the pronunciation of phoneme *k* is dependent on its neighboring phonemes in cat (*k-ae-t*), car (*k-aa-r*), hack (*hh-ae-k*), sky (*s-k-ay*). However, in the sentence *"this is a cat" (dh-ih-s ih-z ah k-ae-t)*, the pronunciation of phoneme *k* is dependent only on neighboring phones, ***ah*** and ***ae*** but is independent of all the other phonemes in the sentence [21]. CNNs have been researched in the field of monolingual ASR as a feature extraction layer to improve the WERs on DNNs while maintaining the computational efficiency by *learning* the short term contextual dependencies on the acoustic frames [12].

In this paper, we build upon our prior work based on a joint acoustic model for low resource languages [1]. Inspiration from the work done on CNNs in [11] and [12] is taken to learn short term contextual temporal features on acoustic frames using convolutional layers. We experiment with both end-to-end and conventional ASR approaches based on CNN-DNN architectures to draw a comparative analysis between the two of them. Three low resource Indic languages, Gujarati, Telugu, and Tamil are used with a combined training dataset of 75 hours. WER is used as the evaluation metric for all the experiments.

The rest of the paper is organized as follows. Section 2 describes multilingual acoustic models. Section 3 describes the various system components of our experiments. Section 4 describes the experimental setup. Section 5 presents results with a comprehensive and comparative analysis. Section 6 concludes our paper by stating the inferences drawn.

2. Multilingual Acoustic Models

Over the years, multilingual acoustic models have been extensively researched for building robust speech recognition sys-

tems. One way to build such a system is to leverage the shared phone space between languages by combining individual phone sets to build a common phone set. The shared phonetic space increases the training data for individual phonemes enabling the model to learn better acoustic representations. This is especially beneficial for low resource languages.

Several techniques have been used to build such multilingual systems. Strategies like shared hidden layers [22], bottleneck features [23], multitask learning [24] train monolingual models on multilingual data. One approach to building such monolingual models is to share the hidden representations between the languages while keeping the output layer language-specific. Such systems are operated by either placing a Language Identification (LID) front-end to switch to the corresponding monolingual model or by selecting the best hypothesis obtained by running all the monolingual models. The performance of the former approach is largely dependent on the robustness and accuracy of the front-end LID system while the latter approach requires multiple monolingual models to be operated parallelly.

Most individuals in multilingual countries like India engage in code-switching during a spoken conversation [25]. Code-switching is a phenomenon of mixing multiple languages in a single utterance. Due to the language dependency in the above-mentioned models, operating such models in a code-switched environment becomes very tricky. To employ large scale speech recognition systems in such multilingual countries, it is necessary to build language-agnostic acoustic models.

This study uses the common phone set to build a computationally efficient language independent joint acoustic model capable of handling multiple languages in a single system. A parser is used to convert utf8 text format to a language-independent IT3 format [26]. The IT3 format text is then used to generate pronunciation sequences for all the words.

3. System Components

3.1. Convolutional Neural Networks

The domain of images and videos often witness very high dimensional data commonly in the range of 10^6. Training deep neural networks on such high dimensional data generates noise while tremendous machine power is needed to process such large fully connected layers. Convolutional neural networks reduce dimensionality whilst retaining important spatial features such as contour boundaries, edges, simple curves, etc. The extracted low dimensional features are then fed to a discriminative model to perform classification tasks.

A convolutional neural network is made up of stacked convolutional layers where each convolutional layer is a combination of convolution, activation, and pooling.

Mathematically, a non-strided convolution operation is expressed as:

$$C[m,n] = \sum_j \sum_l k[j,l]i[m+j, n+l] \qquad (1)$$

where non-strided denotes that the sliding step for the filter is 1. C is the output of the convolution operation, m and n are the convolution output indices, k is a filter, also known as kernel and i is multi-dimensional input. j and l are the dimensions of the filter.

The output dimension of the non-strided convolution operation is given as

$$d(C) = [p - j + 1, q - l + 1] \qquad (2)$$

where p and q are the input dimensions.

Speech signals, like images, are high dimensional data where just a 1-second long speech sample with a frame rate of 16000Hz consists of 16000 features. Each feature represents a sample in time. Conventionally, generic hand-crafted features like MFCC are extracted to reduce the dimensionality of the speech signals and represent the time-variation as a fixed-dimension vector.

CNNs as a replacement for hand-crafted features have been extensively researched [11–15] to extract contextual temporal features.

3.2. Sampling Raw Signals

Raw input speech signals have variable lengths, whereas CNN expects inputs to be of fixed dimension. Hence, the raw signals are first sampled into feature vectors of fixed dimensions. The speech signals in our dataset have a frame rate of 16kHz and are sampled every 10 ms over a window of 25ms.

3.3. Context Window

Sequential models like RNNs capture coarticulation across consecutive phones. Such models are inherently complex and computationally intensive. Moreover, such models are equipped to capture long term contextual dependencies wherein models capturing minimum duration short term contextual dependencies lead to similar or improved results [21].

CNNs can be trained to capture such short-term contextual dependencies by providing context to the current frame. Context windows represent the frames on the left and the right of the current frame. Instead of passing just the current frame along with the phone label to the model, a window of length $(m + n + 1) \times f$ is passed, where m represents the number of context windows on the left, n is the number of context windows on the right and f is the frame length.

4. Experimental Setup

4.1. Data

The Data is a subset of Interspeech 2018's Low Resource Speech Recognition Challenge for Indian languages by Microsoft and SpeechOcean.com dataset [1] [27].

Table 1: *#Utterances included in training, dev and test set*

Languages	Train Set	Dev Set	Test Set
Gujarati	18307	4500	3075
Telugu	12667	3200	3040
Tamil	15712	3900	3081

India is a country with more than 1500 recognized languages. Out of these, 30 languages have more than one million speakers and 22 languages have been accorded with the official status [28]. Such diversity in spoken languages poses a significant challenge in obtaining sizable training data to train robust monolingual systems for each of these languages. This dataset

---

[1]The dataset is available at `https://msropendata.com/datasets/7230b4b1-912d-400e-be58-f84e0512985e`.
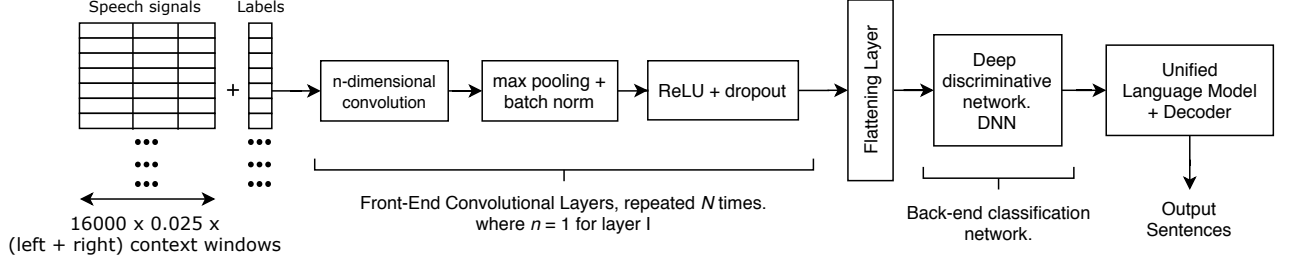
Figure 1: *Design of proposed CNN based architecture*

was released as an effort to explore robust multilingual systems as a means to overcome the challenge of data limitations.

The data includes three Indic languages namely Gujarati, Tamil, and Telugu spoken by multiple speakers. The combined training data of all three languages is 75 hours. The test and the validation data are 5 hours per language. Text transcription along with the lexicon for the entire data is included in the dataset.

State of the art systems were built on the full dataset, which has 120 hours of trainable data. The baseline DNN based system in the challenge has a WER of 27.79, 34.97, and 25.47 on Gujarati, Telugu, and Tamil respectively [27].

4.2. Toolkits

The Kaldi Speech recognition toolkit [29] has been used. SRILM toolkit is used for language modeling. MFCCs, modified KneserNey smoothed tri-gram models are built for forced alignments using the Wall Street Journal Kaldi recipe. LibriSpeech recipes are used to decode and score the system. Pytorch-Kaldi toolkit [30] was used for the development of acoustic models.

4.3. Baseline: Recurrent Neural Network

The baseline model is a recurrent network composed of stacked bidirectional long short-term memory (LSTM) layers built on MFCC features. The network has 3 hidden layers and 550 hidden units in each layer. tanh activation function is used with a 20% dropout in each layer. Layer normalization is used along with a batch size of 8. The learning rate is kept at 0.0016 with a halving factor of 0.5 with the improvement threshold of 0.001. RMSprop optimizer function is used. The MFCC features are generated using the techniques mentioned in section 4.2. Negative Loss Likelihood (NLL) is used as the loss function for the baseline and all the subsequent experiments mentioned in the paper. It is expressed as:

$$L(y) = - \sum_i log(y_i) \qquad (3)$$

where y_i is the prediction for the i^{th} frame.

4.4. Proposed Architecture

Our CNN based acoustic system consists of two components, the front-end and the back end. The front-end is a generative feature extraction network based on deep CNNs and the back end is a discriminative network based on DNNs. The components are connected by a flattening layer and are trained and optimized as a single unit.

Figure 1 represents the high-level design for the proposed end-to-end architecture. Each CNN layer performs convolutions, max pooling, batch normalization with rectified linear units (ReLU) activation. The DNNs consist of stacked fully connected layers with each layer made up of 1024 ReLUs.

The end-to-end system is trained on a feature vector - label pair where each feature vector is a concatenation of left context window, current acoustic frame, and right context window. The utterances are decoded with the decoder and a unified language model mentioned in section 4.2. The flattening layer flattens the multi-channel CNN output as a single-dimensional input to the discriminative network.

5. Experimental Results

This section presents a comprehensive and comparative analysis between different CNN based architectural configurations and the baseline RNN.

Table 2 presents a comprehensive view of the experimental results obtained on models with different configurations. Table 3 presents a comparative analysis of computational time and average WER degradation of different configurations against the baseline model. The unit of the training time is kept variable to enable better readability. The inference time is calculated as T/N where T is the total inference time on the test set and N is the number of examples. All the models are trained, and inferences are drawn on single-core NVIDIA TITAN Xp GPUs in a multi-core GPU setup.

Table 2: *WER of different models in %*

Models + Context	Gujarati	Telugu	Tamil
lstm + mfcc	**18.36**	**25.23**	**24.32**
cnn + raw + $\{0,0\}$	24.06	31.23	30.96
cnn + raw + $\{-1,+1\}$	23.92	30.66	29.58
cnn + raw + $\{-2,+1\}$	**21.13**	**27.80**	**25.76**
cnn + raw + $\{-1,+2\}$	22.65	30.32	28.96
cnn + raw + $\{-2,+2\}$	23.87	32.09	31.20
cnn + mfcc	25.05	31.13	30.78

Two different CNN based architectures are trained. The first architecture is an end-to-end system trained on raw speech signals to enable the CNNs to *learn* sophisticated short-term contextual features directly on the raw data. The second architecture is based on the conventional ASR approach where the CNNs are trained on handcrafted MFCC features.

Table 3: *Performance comparison on the average % WER, average % WER degration, training and inference time.*

Exp	Models + Context	Avg. WER	Avg. WER deg.	Training		Inference	
				Time	Speed Up	Time	Speed Up
1	lstm + mfcc	**22.63**	-	~ 4.5 days	-	780 ms	-
2	cnn + raw + $\{0, 0\}$	28.75	-6.12	7.84 hours	$\sim 13.5\times$	15ms	$\sim 52\times$
3	cnn + raw + $\{-1, +1\}$	28.05	-5.42	11.56 hours	$\sim 9\times$	29ms	$\sim 26\times$
4	cnn + raw + $\{-2, +1\}$	**24.89**	-2.26	19.08 hours	$\sim 5.6\times$	30ms	$\sim 26\times$
5	cnn + raw + $\{-1, +2\}$	27.31	-4.68	21.63 hours	$\sim 5\times$	32ms	$\sim 24\times$
6	cnn + raw + $\{-2, +2\}$	29.05	-6.42	24.42 hours	$\sim 4.4\times$	89ms	$\sim 8.7\times$
7	cnn + mfcc	28.98	-6.35	3.55 hours	$\sim 30\times$	12 ms	$\sim 65\times$

5.1. Deep Convolutional Layers on Raw Speech Signals

The front-end is a CNN composed of 3 hidden layers. The 1[st], 2[nd] and 3[rd] layer has 8, 8 and 2 channels respectively. The length of the kernel in the first layer is 128. Kernel length is halved for each consecutive layer. Max pooling of length 5 is applied on the first layer and a length of 3 is applied on the other two layers. Drop out of 15%, 30% and 20% is employed for the 1[st], 2[nd] and 3[rd] layer respectively. A learning rate of 0.0008 is employed with a halving factor of 0.5 and an improvement threshold of 0.001. ReLU activation function is employed at each layer.

The back end DNN consists of 5 hidden layers. Each of the hidden layers is made up of 1024 units with a dropout of 10%. ReLU activation is applied to each layer with softmax activation on the output layer. A learning rate of 0.0004 is applied with a halving factor of 0.5 and an improvement threshold of 0.001. The end-to-end system is trained on mini-batches of 64 with batch normalization and RMSprop optimizer.

The architecture is trained on raw speech signals sampled to obtain 400-dimensional feature vectors using the technique mentioned in section 3.2. Different experiments are performed varying the size of the context window from 0 to 2 while keeping the CNN and DNN layers constant to find the ideal configuration. Adding 2 context windows on the left and right expands the 400 dimensional feature vector to $(2 + 2 + 1) \times 4 = 2000$ dimension.

Our best model achieves a WER of 21.13, 27.80, and 25.76 on Gujarati, Telugu, and Tamil respectively with 2 left and 1 right context window (Exp 4, Table 3). The results are comparable to the baseline with an average degradation of only 2.26 WER. A computation boost of $5.6\times$ in training time and $26\times$ in inference time is obtained with this configuration taking only an average of 30ms to infer on an average of 5.85s long speech utterance compared to 780ms on the baseline.

It is observed from Table 3 that decreasing the size of the context window boosts up the training and inference computations but at a cost of degradation in the WERs. Models trained on acoustic frames with no context (Exp 2) and large context (Exp 6) observe a significant average degradation of ~ 6.2WER. This indicates no context doesn't provide all the relevant articulation and coarticulation information for the phones while providing large contexts introduces more noise than relevancy. It is also observed that adding more context on the left (Exp 4) gives slightly better WER as compared to adding more context on the right (Exp 5) for the same context window size with comparable training and inference time speed up indicating that the left context contains more relevant information compared to the right context.

5.2. Deep Convolutional Layers on MFCC features

The front-end is a deep CNN composed of 3 hidden layers. The layers consist of 80, 60, 60 output channels respectively. The kernel length of 10, 3, 3 with max-pooling length of 3, 2, 1 is applied. The back end is a deep neural network with 4 hidden layers, each layer composed of 1024 units. A dropout of 15% is applied in each layer of the end to end system with a learning rate of 0.08 and a halving factor of 0.5 with the improvement threshold of 0.001. ReLU activation is applied in every layer barring the output layer which employs softmax activation. The system is trained on a batch size of 128 with batch normalization and stochastic gradient descent optimizer.

The highest boost in the training and inference time is achieved from this configuration (Exp 7, Table 3) with an inference time of an average $12ms$ on an average 5.85s long speech utterance. This is a $65\times$ boost compared to $780ms$ inference time on the baseline. However, this comes with a significant degradation in the WER with an average degradation of 6.35. The WER performance is on par with the CNN architecture on raw speech signals with no context window with a speedup of 3ms on inference time.

6. Conclusion

Our previous work proposed a joint acoustic model based on HMM-SGMM and RNN-CTC for training a multilingual speech recognition system. In this paper, a non-sequential discriminative approach based on the features extracted by CNNs is investigated. Experimental results show that such systems vastly reduce the training and inference time while producing comparable WERs against our baseline sequential model based on RNN. This shows great potential and promise for CNN based architectures in real-time applications for achieving low latency and can also be used to quickly bootstrap multiple low resource multilingual ASR systems. It is observed that a context window of size 2 on the left and 1 on the right produces the most optimal WERs. The best inference time was obtained on the model trained on MFCC features boosting the inference time by $65\times$. CNN based architectures can be further researched to improve and obtain better WERs compared to sequential models while speeding up the training and inference time.

7. Acknowledgements

8. References

[1] H. Krishna, K. Gurugubelli, V. V. R. V, and A. K. Vuppala, "An exploration towards joint acoustic modeling for indian languages: Iiit-h submission for low resource speech recognition challenge for indian languages, interspeech 2018," in *Proc. Interspeech 2018*, 2018, pp. 3192–3196. [Online]. Available: http://dx.doi.org/10.21437/Interspeech.2018-1584

[2] N. Fathima, T. Patel, M. C, and A. Iyengar, "Tdnn-based multilingual speech recognition system for low resource indian languages," 09 2018, pp. 3197–3201.

[3] B. Pulugundla, M. K. Baskar, S. Kesiraju, E. Egorova, M. Karafiát, L. Burget, and J. Černocký, "But system for low resource indian language asr," 09 2018, pp. 3182–3186.

[4] A. Biswas, E. Yılmaz, F. de Wet, E. van der Westhuizen, and T. Niesler, "Semi-supervised acoustic model training for five-lingual code-switched asr," 2019.

[5] A. Kannan, A. Datta, T. N. Sainath, E. Weinstein, B. Ramabhadran, Y. Wu, A. Bapna, Z. Chen, and S. Lee, "Large-scale multilingual speech recognition with a streaming end-to-end model," 2019.

[6] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-to-end continuous speech recognition using attention-based recurrent nn: First results," 2014.

[7] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4960–4964.

[8] T. Hori, S. Watanabe, Y. Zhang, and W. Chan, "Advances in joint ctc-attention based end-to-end speech recognition with a deep cnn encoder and rnn-lm," 2017.

[9] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent 'networks," vol. 2006, 01 2006, pp. 369–376.

[10] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*, ser. ICML'14. JMLR.org, 2014, p. II–1764–II–1772.

[11] J. Li, V. Lavrukhin, B. Ginsburg, R. Leary, O. Kuchaiev, J. M. Cohen, H. Nguyen, and R. T. Gadde, "Jasper: An end-to-end convolutional neural acoustic model," 2019.

[12] P. Golik, Z. Tüske, R. Schlüter, and H. Ney, "Convolutional neural networks for acoustic modeling of raw time signal in lvcsr," in *INTERSPEECH*, 2015.

[13] V. Passricha and R. K. Aggarwal, "Convolutional neural networks for raw speech recognition," in *From Natural to Artificial Intelligence*, R. Lopez-Ruiz, Ed. Rijeka: IntechOpen, 2018, ch. 2. [Online]. Available: https://doi.org/10.5772/intechopen.80026

[14] T. N. Sainath, R. J. Weiss, A. W. Senior, K. W. Wilson, and O. Vinyals, "Learning the speech front-end with raw waveform cldnns," in *INTERSPEECH*, 2015.

[15] D. Palaz, M. Magimai.-Doss, and R. Collobert, "Convolutional neural networks-based continuous speech recognition using raw speech signal," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 4295–4299.

[16] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 1021–1028.

[17] K. Precoda, "Non-mainstream languages and speech recognition: Some challenges," *CALICO Journal*, vol. 21, no. 2, pp. 229–243, January 2004. [Online]. Available: https://www.learntechlib.org/p/74318

[18] A. Graves, A. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 6645–6649.

[19] H. Sak, A. Senior, and F. Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, pp. 338–342, 01 2014.

[20] S. Zhou, Y. Zhao, S. Xu, and B. Xu, "Multilingual recurrent neural networks with residual learning for low-resource speech recognition," in *INTERSPEECH*, 2017.

[21] A. Senior, H. Sak, and I. Shafran, "Context dependent phone models for lstm rnn acoustic modelling," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 4585–4589.

[22] A. Ghoshal, P. Swietojanski, and S. Renals, "Multilingual training of deep neural networks," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7319–7323.

[23] S. Thomas, S. Ganapathy, and H. Hermansky, "Multilingual mlp features for low-resource lvcsr systems," 03 2012, pp. 4269–4272.

[24] D. Chen and B. K. Mak, "Multitask learning of deep neural networks for low-resource speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 7, pp. 1172–1183, 2015.

[25] E. Yılmaz, H. van den Heuvel, and D. Van Leeuwen, "Code-switching detection using multilingual dnns," 12 2016.

[26] A. Baby, N. N. L., A. L. Thomas, and H. A. Murthy, "A unified parser for developing indian language text to speech synthesizers," in *Text, Speech, and Dialogue - 19th International Conference, TSD 2016, Brno, Czech Republic, September 12-16, 2016, Proceedings*, ser. Lecture Notes in Computer Science, P. Sojka, A. Horák, I. Kopecek, and K. Pala, Eds., vol. 9924. Springer, 2016, pp. 514–521. [Online]. Available: https://doi.org/10.1007/978-3-319-45510-5_59

[27] B. Srivastava, S. Sitaram, R. Mehta, K. Mohan, P. Matani, S. Satpal, K. Bali, R. Srikanth, and N. Nayak, "Interspeech 2018 low resource automatic speech recognition challenge for indian languages," 08 2018, pp. 11–14.

[28] Wikipedia contributors, "Language — Wikipedia, the free encyclopedia," 2020, [Online; accessed 12-May-2020]. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Language&oldid=955858850

[29] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The kaldi speech recognition toolkit," in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society, Dec. 2011, iEEE Catalog No.: CFP11SRW-USB.

[30] M. Ravanelli, T. Parcollet, and Y. Bengio, "The pytorch-kaldi speech recognition toolkit," in *In Proc. of ICASSP*, 2019.