



Multi-Modal Data Augmentation for End-to-End ASR

Adithya Renduchintala, Shuoyang Ding, Matthew Wiesner, Shinji Watanabe

Center for Language and Speech Processing
Johns Hopkins University, Baltimore, MD 21218, USA

{adi.r,dings,wiesner,shinjiw}@jhu.edu

Abstract

We present a new end-to-end architecture for automatic speech recognition (ASR) that can be trained using *symbolic* input in addition to the traditional acoustic input. This architecture utilizes two separate encoders: one for acoustic input and another for symbolic input, both sharing the attention and decoder parameters. We call this architecture a multi-modal data augmentation network (MMDA), as it can support multi-modal (acoustic and symbolic) input and enables seamless mixing of large text datasets with significantly smaller transcribed speech corpora during training. We study different ways of transforming large text corpora into a symbolic form suitable for training our MMDA network. Our best MMDA setup obtains small improvements on character error rate (CER), and as much as 7-10% relative word error rate (WER) improvement over a baseline both with and without an external language model.

1. Introduction

The simplicity of “end-to-end” models and their recent success in neural machine-translation (NMT) have prompted considerable research into replacing conventional ASR architectures with a single “end-to-end” model, which trains the acoustic and language models jointly rather than separately. Recently, [1] achieved state-of-the-art results using an attention-based encoder-decoder model trained on over 12K hours of speech data. However, on large publicly available corpora, such as “Librispeech” or “Fisher English”, which are an order of magnitude smaller, performance still lags behind that of conventional systems. [2, 3, 4]. Our goal is to leverage much larger text corpora alongside limited amounts of speech datasets to improve the performance of end-to-end ASR systems.

Various methods of leveraging these text corpora have improved end-to-end ASR performance. [5], for instance, composes RNN-output lattices with a lexicon and word-level language model, while [6] simply re-scores beams with an external language model. [7, 8] incorporate a character level language model during beam search, possibly disallowing character sequences absent from a dictionary, while [9] includes a full word level language model in decoding by simultaneously keeping track of word histories and word prefixes. As our approach does not change any aspect of the traditional decoding process in end-to-end ASR, the methods mentioned above can still be used in conjunction with our MMDA network.

An alternative method, proposed for NMT, augments the source (input) with “synthetic” data obtained via back-translation from monolingual target-side data [10]. We draw inspiration from this approach and attempt to augment the ASR input with text-based synthesized input generated from large text corpora.




Figure 1: Overview of our Multi-modal Data Augmentation (MMDA) model. Figure 1a highlights the network engaged when acoustic features are given as input to an acoustic encoder (shaded blue). Alternatively, when synthetic input is supplied the network (Figure 1b) uses an augmenting encoder (green). In both cases a shared attention mechanism and decoder are used to predict the output sequence. For simplicity we show 2 layers without down-sampling in the acoustic encoder and omit the input embedding layer in the augmenting encoder.

2. Approach

While text-based augmenting input data is a natural fit for NMT, it cannot be directly used in end-to-end ASR systems which expect acoustic input. To utilize text-based input, we use two separate encoders in our ASR architecture: one for acoustic input and another for synthetic text-based augmenting input. Figure 1 gives an overview of our proposed architecture.

2.1. MMDA Architecture

Figure 1a shows a sequence of acoustic frames $\{x_0, x_1, \dots\}$ fed into an acoustic encoder shown with blue cross hatching. The attention mechanism takes the output of the encoder and generates a context vector (gray cross hatching) which is uti-

Table 1: Examples of sequences under different synthetic input generation schemes. The original text for these examples is the phrase JOHN BLARE AND COMPANY.

Synthetic Input	Example Sequence
Charstream	J O H N B L A R E A N D C O M P A N Y
Phonestream	JH AA1 N B L EH1 R AE1 N D K
	AH1 M P AHO N IYO
Rep-Phonestream	JH JH JH AA1 AA1 AA1 AA1 N
	B L L L EH1 R AE1 AE1 AE1 N D K K K K AH1
	AH1 M M P AHO AHO AHO N IYO IYO IYO

lized by the decoder (red cross hatching) to generate each token in the output sequence $\{y_0, y_1, \dots\}$. In figure 1b, the network is given a sequence of “synthetic” input tokens, $\{z_0, z_1, \dots\}$, where $z_i \in \mathcal{Z}$ and the set \mathcal{Z} is the vocabulary of the synthetic input. The size and items in \mathcal{Z} depend on the type of synthetic input scheme used (see Table 1 for examples and Section 5.2 for more details). As the synthetic inputs are categorical, we use an input embedding layer which learns a vector representation of each symbol in \mathcal{Z} . The vector representation is then fed into an augmenting encoder (shown in green cross hatching). Following this, the *same* attention mechanism and decoder are used to generate an output sequence. Note that some details such as the exact number of layers, down-sampling in the acoustic encoder, and the embedding layer in the augmenting encoder are omitted in Figure 1 for sake of clarity.

2.2. Synthetic Inputs

A desirable synthetic input should be easy to construct from plain text corpora, and should be as similar as possible to acoustic input. We propose three types of synthetic inputs that can be easily generated from text corpora and with varied similarity to acoustic inputs (see Table 1).

1. *Charstream*: The output character sequence is supplied as synthetic input without word boundaries.
2. *Phonestream*: We make use of a pronunciation lexicon to expand words into phonemes where unknown pronunciations are recovered via grapheme-to-phoneme transduction (G2P).
3. *Rep-Phonestream*: We explicitly model phoneme duration by repeating each phoneme such that the relative durations of phonemes to each other mimic what is observed in data (e.g. vowels last longer than stop consonants).

2.3. Multi-task Training

Let \mathcal{D} be the ASR dataset, with acoustic input and character sequence output pairs $(\mathbf{X}_j, \mathbf{y}_j)$ where $j \in \{1, \dots, |\mathcal{D}|\}$. Using a text corpus \mathcal{S} with sentences \mathbf{s}_k where $k \in \{1, \dots, |\mathcal{S}|\}$, we can generate synthetic inputs $\mathbf{z}_k = \text{syn}(\mathbf{s}_k)$, where $\text{syn}(\cdot)$ is one of the synthetic input creation schemes discussed in the previous section. Under the assumption that both \mathbf{y}_j and \mathbf{s}_k are sequences with the same character vocabulary and from the same language, our augmenting dataset \mathcal{A} is comprised of training pairs $(\mathbf{z}_k, \mathbf{s}_k)$ $k \in \{1, \dots, |\mathcal{S}|\}$. Typically the corpus \mathcal{S} is much larger than the original ASR training set \mathcal{D} . During training, we alternate between batches from acoustic training data \mathcal{D} (primary task) and synthesized augmenting data \mathcal{A} (secondary

task). In each batch we maximize the primary objective or the secondary objective. Note that in both cases the attention and decoder parameters (denoted by θ_{att} and θ_{dec} , see equation 1) are shared, while the acoustic encoder parameters (θ_{enc}) and augmenting encoder parameters (θ_{aug}) are only updated in their respective training batches.

$$\mathcal{L}(\theta) = \begin{cases} \log P(\mathbf{y} | \mathbf{X}; \theta_{\text{enc}}, \theta_{\text{att}}, \theta_{\text{dec}}) & \text{primary objective} \\ \log P(\mathbf{s} | \mathbf{z}; \theta_{\text{aug}}, \theta_{\text{att}}, \theta_{\text{dec}}) & \text{secondary objective} \end{cases} \quad (1)$$

We evaluate our model on a held out ASR dataset \mathcal{D}' which only contains acoustic batches as our ultimate goal is to obtain the best ASR system.

In the remainder of the paper we place our work in context of other multi-modal, multi-task, and data-augmentation schemes for ASR. We propose a novel architecture to seamlessly train on both text (with synthetic inputs) and speech corpora. We analyze the merit of these approaches on WSJ, and finally report the performance of our best performing architecture on WSJ [11] and the HUB4 Spanish [12] and Voxforge Italian corpora [13].

3. Related Work

Augmenting the ASR source with synthetically generated data is already a widely used technique. Generally, label-preserving perturbations are applied to the ASR source to ensure that the system is robust to variations in source-side data not seen in training. Such perturbations include Vocal Tract Length Perturbations (VTLN) as in [14] to expose the ASR to a variety of synthetic speaker variations, as well as speed, tempo and volume perturbations [15]. Speech is also commonly corrupted with synthetic noise or reverberation [16, 17].

Importantly, these perturbations are added to help learn more robust acoustic representations, but not to expose the ASR system to new output utterances, nor do they alter the network architecture. By contrast, our proposed method for data augmentation from external text exposes the ASR system to new output utterances, rather than to new acoustic inputs.

Another line of work involves data-augmentation for NMT. In [18], improvements in low-resource settings were obtained by simply copying the source-side (input) monolingual data to the target side (output). Our approach is loosely based off of [10], which improves NMT performance by creating pseudo parallel data using an auxiliary translation model in the reverse direction on target-side text.

Previous work has also tried to incorporate other modalities during both training and testing, but have focused primarily on learning better feature representations via correlative objective functions or on fusing representations across modalities [19, 20]. The fusion methods require both modalities to be present at test time, while the multiview methods require both views to be conditionally independent given a common source. Our method has no such requirements and only makes use of the alternate modality during training.

Lastly, we note that considerable work has applied multi-task training to “end-to-end” ASR. In [21], the CTC objective is used as an auxiliary task to force the attention to learn monotonic alignments between input and output. In [22], a multi-task framework is used to jointly perform language-id and speech-to-text in a multilingual ASR setting. In this work our use of phoneme-based augmenting data is effectively using G2P (P2G) as an auxiliary task in end-to-end ASR, though only implicitly.

4. Method

Our MMDA architecture is a straightforward extension to Attention-Based Encoder-Decoder network [23], which is described by components as follows.

4.1. Acoustic Encoder

For a single utterance, the acoustic frames form a matrix $\mathbf{X} \in \mathbb{R}^{L_x \times D_x}$ are encoded by a multi-layer bi-directional LSTM (biLSTM) with hidden dimension H for each direction, L_x and D_x being the length of the utterance in frames and the number of acoustic features per frame, respectively. After each layers' encoding, the hidden vectors of \mathbb{R}^{2H} are projected back to vectors of \mathbb{R}^H using a projection layer and fed as the input into the next layer. We also use a pyramidal encoder following [24] to down-sample the frame encodings and capture a coarser-grained resolution.

4.2. Augmenting Encoder

The augmenting encoder is a single-layer biLSTM – essentially a “shallow” acoustic encoder. As the synthetic input is symbolic (e.g. phoneme, character), we use an embedding layer which learns a real-valued vector representation for each symbol, thus converting a sequence of symbols $\mathbf{z} \in \mathcal{Z}^{L_z \times 1}$ into a matrix $\mathbf{Z} \in \mathbb{R}^{L_z \times D_z}$, where \mathcal{Z} is the set of possible augmenting input symbols, L_z is the length of the augmenting input sequence and D_z is embedding size respectively. We set $D_x = D_z$ to ensure that the acoustic and augmenting encoders work smoothly with the attention mechanism.

4.3. Decoder

We used a uni-directional LSTM for the decoder [23, 25].

$$\mathbf{s}_j = \text{LSTM}(\mathbf{y}_{j-1}, \mathbf{s}_{j-1}, \mathbf{c}_j) \quad (2)$$

where \mathbf{y}_{j-1} is the embedding of the last output token, \mathbf{s}_{j-1} is the LSTM hidden state in the previous time step, and \mathbf{c}_j is the attention-based context vector which will be discussed in the following section. We omitted all the layer index notations for simplicity. The hidden state of the final LSTM layer is passed through another linear transformation followed by a softmax layer generating a probability distribution over the outputs.

4.4. Attention Mechanism

We used Location-aware attention [26], which extends the content-based attention mechanism [23] by using the attention weights from the previous output time-step α_{j-1} when computing the weights for the current output α_j . The previous time-step attention weights α_{j-1} are “smoothed” by a convolution operation and fed into the attention weight computation. Once attention weights are computed, a weighted sum over encoder hidden states generates the context vector \mathbf{c}_j .

5. Experiments

5.1. Data

We compared the proposed types of synthetic data by evaluating character and word error rates (CER, WER) of ASR systems trained on the Wall Street Journal corpus (LDC93S6B and LDC94S13B), using the standard SI-284 set containing $\sim 37\text{K}$ utterances or 80 hours of speech. We used the “dev93” set as a development set and selection criteria for the best model, which

was then evaluated on the “eval92” dataset. We also tested the performance of MMDA using the best performing synthetic input type on the Hub4 Spanish and Italian Voxforge datasets.

The Hub4 Spanish corpus consists of 30 hours of 16kHz speech from three different broadcast news sources [12]. We used the same evaluation set as used in the Kaldi Hub4 Spanish recipe [27], and constructed a development set with the same number of utterances as the evaluation set by randomly selecting from the remaining training data.

For the Voxforge Italian corpus, which consist of 16 hours of broadband speech, [13] we created training, development, and evaluation sets, by randomly selecting 80%, 10%, and 10% of the data for each set respectively, ensuring that no sentence was repeated in any of the sets. In all experiments we represented each frame of audio by a vector of 83 dimensions (80 Mel-filter bank coefficients 3 pitch features).

5.2. Generating Synthetic Input

The augmenting data used for the WSJ experiments are generated from section (13-32.1 87,88,89) of WSJ, which is typically used for training language models applied during decoding. We made 3 different synthetic inputs for this section of WSJ. For *Charstream* synthetic input the target-side character sequence was copied to the input while omitting word boundaries. For *Phonestream* synthetic input we constructed phone sequences using CMUDICT to which 46k words from the WSJ corpus are added[28] as the lexicon as described in section 5.2. We trained the G2P on CMUDICT using the *Phonetisaurus* toolkit. For certain words consisting only of rare graphemes, we were unable to infer pronunciations and simply assigned to these words a single ⟨unk⟩ phoneme. Finally, we filtered out sentences with more than 1 ⟨unk⟩ phoneme symbol, and those above 250 characters in length. The resulting augmenting dataset contained $\sim 1.5M$ sentences.

In the *Rep-Phonestream* scheme, we modified the augmenting input phonemes to further emulate the ASR input by modeling the variable durations of phonemes. We assumed that a phoneme’s duration in frames is normally distributed $\mathcal{N}(\mu_p, \sigma_p^2)$ and we estimated these distributions for each phoneme from frame-level phoneme transcripts in the TIMIT dataset. For example, given a phoneme sequence like JH AA N (for the word “John”), we would sample a sequence of frame durations $f_p \sim \mathcal{N}(\mu_p, \sigma_p^2)$ $p \in \{\text{JH}, \text{AA}, \text{N}\}$ and repeat each phoneme r times, where $r = \max(1, \text{Round}(f_p)/4)$. Dividing by 4 accounts for the down-sampling performed by the pyramidal scheme in the acoustic encoder.

The augmenting data for both Spanish and Italian was generated by using Wikipedia data dumps¹ and then scraping Wiktionary using *wikt2pron*² for pronunciations of all words seen in the text. We used the resulting seed pronunciation lexicon for G2P training as before and again filtered out long sentences and those with resulting ⟨unk⟩ words after phonemic expansion of all words in the augmenting data. In order to generate the *Rep-Phonestream* data we manually mapped TIMIT phonemes to similar Italian and Spanish phonemes and applied the corresponding durations learned on TIMIT.

5.3. Training

We implemented our MMDA model on-top of ESPNET using the PyTorch backend [21, 29]. A 4 layer biLSTM with a “pyra-

¹<https://dumps.wikimedia.org/backup-index.html>

²<https://github.com/abuccts/wikt2pron>

Table 2: Experiments on WSJ corpus using different augmentation input types (part 1). The best performing augmentation was then applied to Italian (Voxforge) and Spanish (HUB4) datasets (part 2 & 3 of the table).

Corpus	Augmentation	CER (eval, dev)	WER (eval, dev)
English (WSJ)	No-Augmentation	7.0 , 9.9	19.5, 24.8
	Charstream	7.5, 10.5	20.3, 25.7
	Phonestream	7.4, 10.1	20.4, 25.3
	Rep-Phonestream	7.1, 9.8	17.5 , 22.7
	No-Augmentation + LM	7.0, 9.8	17.2, 22.2
	Rep-Phonestream + LM	6.7 , 9.4	16.0 , 20.8
Italian (Voxforge)	No-Augmentation + LM	16.4, 15.9	47.2, 46.1
	Rep-Phonestream + LM	14.8 , 14.5	44.3 , 44.0
Spanish (HUB4)	No-Augmentation + LM	12.6, 12.8	31.5, 33.5
	Rep-Phonestream + LM	12.1 , 13.1	29.5 , 32.6

midal” structure was used for the acoustic encoder [6]. The biLSTMs in the encoder used 320 hidden units (in each direction) followed by a projection layer. For the augmenting encoder, we used a single layer biLSTM with the same number of units and projection scheme as the acoustic encoder. No down-sampling was done on the augmenting input. Location-aware attention was used in all our experiments [26]. For WSJ experiments, the decoder was a 2-layer LSTM with 300 hidden units, while a single layer was used for both Spanish and Italian. We used Adadelta to optimize all our models for 15 epochs [30]. The model with best validation accuracy (at the end of each epoch) was used for evaluation.

For decoding, a beam-size of 10 for WSJ and 20 for Spanish and Italian was used. In both cases we restricted the output using a minimum-length and maximum-length threshold. The min and max output lengths were set as $0.3F$ and $0.8F$, where F denotes the length of down-sampled input. For RNNLM integration, we trained a 2-layer LSTM language model with 650 hidden units. The RNNLM for each experiment was trained on the same sentences used for augmentation.

5.4. Results

Table 2 (part 1) shows the ASR results on WSJ. Rep-Phonestream augmentation improved the baseline WER by a margin of 2%, while none of the other augmentations helped. This corroborates our intuition that data augmentation works better when synthetic inputs are similar to the real training data. Furthermore, we continued to observe gains in WER when an RNNLM was incorporated in the decoding process [8]. This suggests that while MMDA and LM have a similar effect, they can still be used in conjunction to extract further improvement.

The best performing synthetic input scheme was applied to Spanish and Italian, where a similar trend was observed. MMDA consistently achieved better WER and obtained small improvements in CER (see table 2, parts 2 and 3). The relative gains in English (WSJ) were higher than Spanish and Italian; we suspect the ad-hoc phone duration mapping we employed for these languages and mismatch in augmenting text data might have contributed to the lower relative gains.

We found that the Rep-Phonestream MMDA system tended to replace entire words when incorrect, while the baseline system incorrectly changed a few characters in a word, even if the resulting word did not exist in English (for WSJ). This behav-

Table 3: Error type differences between the Rep-Phonestream MMDA trained system and the baseline system on WSJ (dev & test combined). “Nonsense errors” are substitutions or insertions that result in non-legal English words, e.g. CASINO substituted with ACCINO. “Legal errors” are errors that result legal English words, e.g. BOEING substituted with BOLDING.

Augmentation	Nonsense errors %	Legal errors %
No-Augmentation	32.93	67.07
No-Augmentation + LM	24.34	75.66
Rep-Phonestream	24.99	75.11
Rep-Phonestream + LM	20.25	79.75

ior tended to improve WER while harming CER. For example, in the WSJ experiments, the baseline substitutes QUOTA with COLOTA, while the Rep-Phonestream MMDA predicts COLORS. We verified this hypothesis by computing the ratio of substitutions and insertions resulting in nonsense words to the total number of such errors on the WSJ development and evaluation data for the baseline system and MMDA, both with and without RNNLM re-scoring. We see that RNNLM re-scoring actually behaves like MMDA in this regard (see Table 3).

6. Future Work

6.1. Enhancing MMDA

We identify three possible future research directions:

- (i) *Augmenting encoders*: More elaborate designs for the augmenting encoder could be used to generate more “speech-like” encodings from symbolic synthetic inputs.
- (ii) *Synthetic inputs*: Other synthetic inputs should be explored, as our choice was motivated in large part by simplicity and speed of generating synthetic inputs. Using something approaching a text-to-speech system to generate augmenting data may be greatly beneficial.
- (iii) *Training schedules*: Is the 1:1 ratio for augmenting and acoustic training ideal? Using more augmenting data initially may be beneficial, and a systematic study of various training schedules would reveal more insights. Furthermore, automatically adjusting the amount of augmenting data used also seems worthy of inquiry.

6.2. Applications

Our framework is easily expandable to other end-to-end sequence transduction applications, examples of which include domain adaptation and speech-translation. To adapt ASR to a new domain (or even new dialect/language), we can train on additional augmenting data derived from the new domain (or dialect/language). We also believe the MMDA framework may be well suited to speech-translation due to its similarity to back-translation in NMT.

7. Conclusion

We proposed the MMDA framework which exposes our end-to-end ASR system to a much wider range of training data. To the best of our knowledge, this the first attempt in truly end-to-end multi-modal data augmentation for ASR. Experiments show promising results for our MMDA architecture and we highlight possible extensions and future research in this area.

8. References

- [1] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, K. Gonina *et al.*, “State-of-the-art speech recognition with sequence-to-sequence models,” *arXiv preprint arXiv:1712.01769*, 2017.
- [2] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015.
- [3] C. Cieri, D. Graff, O. Kimball, D. Miller, and K. Walker, “Fisher english training speech part 1 transcripts,” *Philadelphia: Linguistic Data Consortium*, 2004.
- [4] —, “Fisher english training part 2,” *Linguistic Data Consortium, Philadelphia*, 2005.
- [5] Y. Miao, M. Gowayyed, and F. Metze, “Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding,” in *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*. IEEE, 2015, pp. 167–174.
- [6] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [7] A. Maas, Z. Xie, D. Jurafsky, and A. Ng, “Lexicon-free conversational speech recognition with neural networks,” in *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2015, pp. 345–354.
- [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,” in *Interspeech*, 2017, pp. 949–953.
- [9] A. Graves and N. Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *International Conference on Machine Learning (ICML)*, 2014, pp. 1764–1772.
- [10] R. Sennrich, B. Haddow, and A. Birch, “Improving Neural Machine Translation Models with Monolingual Data,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, August 2016, pp. 86–96. [Online]. Available: <http://www.aclweb.org/anthology/P16-1009.pdf>
- [11] D. B. Paul and J. M. Baker, “The design for the Wall Street Journal-based CSR corpus,” in *Proceedings of the workshop on Speech and Natural Language*. Association for Computational Linguistics, 1992, pp. 357–362.
- [12] L. D. Consortium, “1997 spanish broadcast news speech (hub 4-ne) ldc98s74,” *Philadelphia*, 1998.
- [13] Voxforge.org, “Free speech recognition,” <http://www.voxforge.org/>, accessed 06/25/2014.
- [14] A. Ragni, K. M. Knill, S. P. Rath, and M. J. Gales, “Data augmentation for low resource languages,” in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [15] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [16] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates *et al.*, “Deep speech: Scaling up end-to-end speech recognition,” *arXiv preprint arXiv:1412.5567*, 2014.
- [17] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, “An analysis of environment, microphone and data simulation mismatches in robust speech recognition,” *Computer Speech & Language*, vol. 46, pp. 535–557, 2017.
- [18] A. Currey, A. V. M. Barone, and K. Heafield, “Copied monolingual data improves low-resource neural machine translation,” in *Proceedings of the Second Conference on Machine Translation*, 2017, pp. 148–156.
- [19] R. Arora and K. Livescu, “Multi-view cca-based acoustic features for phonetic recognition across speakers and domains,” in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 7135–7139.
- [20] Y. Mroueh, E. Marcheret, and V. Goel, “Deep multimodal learning for audio-visual speech recognition,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 2130–2134.
- [21] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, “Hybrid CTC/attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [22] S. Toshniwal, T. N. Sainath, R. J. Weiss, B. Li, P. Moreno, E. Weinstein, and K. Rao, “Multilingual speech recognition with a single end-to-end model,” *arXiv preprint arXiv:1711.01694*, 2017.
- [23] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [24] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 4960–4964.
- [25] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, “End-to-end continuous speech recognition using attention-based recurrent nn: First results,” *arXiv preprint arXiv:1412.1602*, 2014.
- [26] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Advances in Neural Information Processing Systems (NIPS)*, 2015, pp. 577–585.
- [27] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, “The kaldi speech recognition toolkit,” in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. EPFL-CONF-192584. IEEE Signal Processing Society, 2011.
- [28] J. Kominek and A. W. Black, “The cmu arctic speech databases,” in *Fifth ISCA Workshop on Speech Synthesis*, 2004.
- [29] S. Kim, T. Hori, and S. Watanabe, “Joint ctc-attention based end-to-end speech recognition using multi-task learning,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 4835–4839.
- [30] M. D. Zeiler, “Adadelta: an adaptive learning rate method,” *arXiv preprint arXiv:1212.5701*, 2012.