END-TO-END MULTIMODAL SPEECH RECOGNITION

Shruti Palaskar*, Ramon Sanabria* and Florian Metze

Carnegie Mellon University; Pittsburgh, PA; U.S.A. {spalaska|ramons|fmetze}@cs.cmu.edu

ABSTRACT

Transcription or sub-titling of open-domain videos is still a challenging domain for Automatic Speech Recognition (ASR) due to the data's challenging acoustics, variable signal processing and the essentially unrestricted domain of the data. In previous work, we have shown that the visual channel – specifically object and scene features – can help to adapt the acoustic model (AM) and language model (LM) of a recognizer, and we are now expanding this work to end-to-end approaches. In the case of a Connectionist Temporal Classification (CTC)-based approach, we retain the separation of AM and LM, while for a sequence-to-sequence (S2S) approach, both information sources are adapted together, in a single model. This paper also analyzes the behavior of CTC and S2S models on noisy video data (How-To corpus), and compares it to results on the clean Wall Street Journal (WSJ) corpus, providing insight into the robustness of both approaches.

Index Terms— Audiovisual Speech Recognition, Connectionist Temporal Classification, Sequence-to-Sequence Model, Adaptation

1. INTRODUCTION

Audio-visual speech recognition has been an active area of research for a long time: humans use "lip-reading" to improve their robustness against noise, and they control the balance between lip-reading and hearing adaptively and transparently.

In this paper, we propose to use multi-modal video information slightly differently, and adapt an end-to-end speech recognizer to the visual semantic concepts extracted from a *correlated visual scene* that accompanies some speech, for example in a "How-To" video. If we see a person standing in a kitchen, holding sliced bread, it is likely that the person is explaining how to make a sandwich, and the acoustic conditions will be comparably clean. If a person is standing in front of a car, it is likely a review of that car, and happening outdoors. Clearly, the language model will be affected as well.

Recognition thus involves two major steps:

- For every utterance, extract a visual semantic feature "context vector" from a single video frame using deep Convolutional Neural Networks (CNNs) trained for object recognition and scene labeling tasks.
- Adapt or condition a recognizer to this utterance's context. In this work, we compare a CTC-based system and a S2S system.

This paper builds on earlier work [1, 2] and presents first results on multi-modal adaptation of CTC and S2S models. We present various audio-visual adaptation strategies for the CTC model and present our first results with S2S model adaptation. The ultimate goal of this work will be to view automatic speech recognition not primarily as the speech-to-text task, but as a process which sub-titles multi-media material removing repetitions, hesitations or corrections from spontaneous speech as required, much like "video captioning" [3]. We show that multi-modal adaptation helps by 2% absolute improvement in the token error rate.

While multi-modal adaptation improves recognition in such noisy datasets, we see that there is need for deeper insight into the CTC and S2S models. These models behave very differently with clean, prepared datasets like WSJ than with spontaneous, noisy speech. Ground truth references for the How-To data are less accurate than for WSJ; we see that this influences model training. We present insights into the differences in the output of these two approaches as these issues have not been addressed in prior work yet. Because of the novelty of the S2S approach, we also implement a strong baseline on the WSJ dataset.

In the following sections, we first describe related work in 2, data and feature extraction in 3, model descriptions in 4 and results on multimodal adaptation in 5.1, differences in CTC and S2S in 5.2 and model differences while using WSJ and How-To in 5.3.

2. RELATED WORK

On consumer-generated content (like YouTube videos), Deep Neural Network (DNN) models exhibit Word Error Rates (WER) above 40% [4], although no standardized test set exists for such type of data. An effective strategy to deal with variability was to incorporate additional, longer-term knowledge explicitly into DNN models: [5, 6, 7, 8] study the incorporation of speaker-level i-vectors to balance the effect of speaker variability. Time Delay Neural Networks [9, 10] use wide temporal input windows to improve robustness. [11] extracts long-term averages from the audio signal to adapt a DNN acoustic model. Similarly, in [12], we learn a DNN-based extractor to model the speaker-microphone distance information dynamically on the frame level. Then distance-aware DNNs are built by appending these descriptors to the DNN inputs. We also try certain speaker and visual adaptive training strategies that result in about 20% WER on the How-To corpus [1, 2]. We follow the trend of including additional, longer-term knowledge but with end-to-end models rather than DNNs.

It is an important distinction that our work *does not require* localization of lip regions and/or extraction of frame-synchronous visual features (lip contours, mouth shape, landmarks, etc.), as is the case in "traditional" audio-visual ASR [13, 14, 15, 16], which has been developed mostly with a focus on noise robustness. For the majority of our data, lip-related information is not available at all, or the quality is extremely poor. Instead, we will use semantic visual features such as objects [17] and scenes [18].

Unlike Hidden Markov Models, end-to-end systems are directly optimized for a sequence of characters, phones, or other target units.

^{*}The first two authors contributed equally to this work.



Fig. 1. Length Distribution for How-To and WSJ Train set

The CTC [19] loss function is defined over a target label sequence, and introduces an additional "blank" label, which the network can predict at any frame without influencing the output label sequence. In CTC training, the sequence of labels are monotonically mapped to the observations (*i.e.*, speech frames), and outputs appear timealigned to inputs. Decoding of CTC models can be achieved with weighted Finite State Transducers (WFSTs, [20, 21]) or a Recurrent Neural Network language model [22]. In the latter case, characters are frequently used as target labels [23]. Finally, attention-based S2S models [24] can be applied to speech recognition as well [25, 26], and present an intriguing alternative to conventional models. Still, their training and decoding are not well understood at this point, in particular on data-sets for which high-quality annotations may not be available. We present an analysis of results with these models to better understand them.

3. DATA AND FEATURES

We conduct our experiments on two data-sets, the Wall Street Journal (WSJ, SI-284, LDC93S6B and LDC94S13B), and the How-To audio-visual dataset. The How-To corpus consists of English language open-domain instructional videos that explain specific tasks like baking a cake, or nutrition habits, and have been recorded in various environments indoors and outdoors (like kitchen, or garden), usually with a portable video recorder [1, 27]. Ground truth transcriptions of these videos have been created by re-aligning provided sub-titles, which sometimes are mis-matched because of missed phrases, word repetitions, hesitations, and other noise and punctuation that hasn't been transcribed.

We use 90 and 480 hours of How-To data and 87 hours of WSJ, and extract 40-dimensional MEL filter banks, with a step size of 30 ms, 3-fold oversampling of the data at 0, 10, and 20 ms offsets, and stacking 3 neighboring frames together, to give the same 120dimensional input vector to both CTC and S2S models. The 90 h subset of How-To has been selected randomly. We have a separate 4 h test set. 5% of the training data is used as dev set. For WSJ, we use the eval92 test set. Both models are character-based with 43 labels/tokens: 26 alphabets, 10 digits, and special symbols for {`.', ``, `-', '/}, space, start and end of sentence.

Figure 1 shows the length distributions for How-To and WSJ train sets. This shows that for "open-domain" speech data, the distribution is less normalized when compared with prepared datasets.

3.1. Extraction of Visual Vectors

The visual features used in this paper are the same as those in our previous work [2]. We extract object and place/scene features from pre-trained CNNs and perform dimensionality reduction to obtain 100 dimension features. As described above, the data contain indoor and outdoor recordings of instructional videos where object and place features are most relevant. We use these features to infer acoustic and language information from the scene where the utterance has been recorded.

4. MODELS

4.1. Audio-Visual ASR using CTC

The CTC acoustic model (AM) is a stacked bidirectional Long Short Term Memory (LSTM) model with a soft-max layer at the end to generate probability of label $k \in L'$ at a particular time step t given the speech sequence **X**. Here, $L' = L \cup \phi$ where L represents the vocabulary described in Section 3 and ϕ denotes the special blank symbol introduced by CTC. To define the CTC loss function we need a many-to-one mapping \mathcal{B} that maps path $\mathbf{p} = (p_1, ..., p_T) \in L'^T$ of the CTC model over length T to an output sequence **z**. The probability of **z** is calculated by summing over all possible paths in CTC.

CTC AM is a frame-based sequence criterion where for every frame, a conditionally independent probability distribution over all labels L' is generated. From this definition, we can infer that CTC is more dependent on the actual acoustics of the input, and can function even with poorly transcribed data. We support this hypothesis with analysis in Section 5.

Decoding strategies to obtain labels from the probability distributions are: greedy search, WFST and character-Recurrent Neural Network (char-RNN). Greedy decoding outputs the character with the highest probability value without taking into account any previous information. We note that this strategy produces good results with the CTC AM and with visual adaptation as well. The second strategy for decoding is to re-score the output of the AM with a wordbased WFST LM like our previous work [21]. Char-RNN is a neural language model like in [22] and provides an all neural model that can be trained end-to-end.

We implement Visual Adaptive Training (VAT) based on our previous work [12, 2]. We jointly train a CTC AM and a separate multi-layer perceptron that is connected to the input LSTM units by a 'sum' operation. This is shown in Figure 2a. The intuition behind this architecture is that the VAT module will perform a normalization over the input features (audio) based on the visual information. With this adaptation, a 'weighted sum' of audio and video features gives us a 120 d features for each frame. Other context information like i-vectors, can be used with this technique. Although in our previous work we trained both models in "two-step" fashion (first AM, then VAT), we experienced higher boost of performance by training both models jointly, end-to-end. Other types of adaptation such as concatenation were tested but none of them outperformed the method discussed.

Our CTC AM uses 5 layers with 200 LSTM cells each, with a projection layer of 100 cells between each LSTM layer. We tried many different architectures to find this best configuration. We use Stochastic Gradient Descent (SGD) for training. Training on 480 h How-To data takes 30 hours and greedy decoding takes less than 15 minutes on TitanX GPU with 11GB RAM.



Fig. 2. Model Architectures

4.2. Audio-Visual ASR using S2S

We implement an attention-based S2S model that is similar to [28, 26]. The encoder transforms the input to a high level, low dimension encoding, $h = (h_1, ..., h_T)$ of length T, and the attention decoder produces a character probability distribution conditioned on all previous output. For each layer in the encoder, we concatenate two consecutive nodes to get a reduction factor of 2 at each layer and form a pyramidal encoder instead of a stacked one. This encoder acts as the 'AM' in CTC. The attention mechanism we implement learns a weighted global context vector W_{α} calculated using the source hidden state h_s and the *current* target h_t [29]. This context vector is global as it always attends to all source states s'. We compute a variable-length alignment vector α_t using: $exp(h_t^T.W_{\alpha}.h_s)$ and this is normalized over all input s' as $\sum_{s'} exp(h_t^T.W_{\alpha}.h_{s'})$. The model architecture with attention is shown in the Figure 2b.

We note that this type of encoder state concatenation and attention mechanism are the novelty of the S2S model in this paper and have not been applied for ASR before. The concatenation method worked slightly better in our experiments than other reduction methods reported in [25, 26, 30]. In addition, a dense-nonlinear layer is used in between the encoder and decoder states. The S2S decoder implicitly learns an LM using the transcripts although this is a weak LM. Decoding is performed using beam search.

S2S is a sequence learning criterion that does not try to output a token for each frame of input. It "attends" over the input and output sequences to align them (multiple frames can be mapped to one token). This leads us to the hypothesis that S2S is much more dependent on the transcripts rather than the acoustics. Again, we support this hypothesis with analysis in Section 5.

In our S2S model, we use 3 layers of 512 bidirectional LSTM cells in the encoder. We use SGD with learning rate of 0.2 and decay of 0.9. We use curriculum learning [31] for the first epoch to speed up convergence. We note that our training process is much simpler than [26, 32, 25]. The decoder is made of 2 layers of 512 bidirectional LSTM cells each. For decoding, we use a beam size of 5. We do not use any techniques for better decoding with WSJ as given in [32] but use a length normalization with How-To data, to address the length distributions variance shown in Figure 1. Training took 4 days with a TitanX GPU on the 90 h subset. Experiments were performed using the OpenNMT toolkit [33].

The video adaptation technique we use with S2S is early fusion

where 100 d visual features are concatenated with 120 d audio features giving 220 d vector for each frame. Our experiments with early fusion show that S2S benefits with this technique while CTC or DNN [12] does not, moreover gives better gains that VAT for CTC.

5. EXPERIMENTS AND ANALYSIS

5.1. Visual Feature Adaptation

In Table 1, we present the effect of adaptation on the Token Error Rate (TER) and Perplexity (PPL) of the two models. We show that the adaptation with visual features helps improve the absolute TERs by 1% in CTC and by 1.6% in the S2S model. Using length norm as in Section 5.3, we improve the S2S improvement to 2%. In our experience, this is a significant improvement in such sequence-based models. The PPL values for CTC are those of a word LM while those of S2S are of the implicit character LM (*). We see that adapting a language model with visual features helps decrease the perplexity by a huge margin. This establishes that there is a strong correlation between visual features and speech. The perplexity for S2S models is the character-prediction perplexity (joint AM, LM) and we see that there is no difference in this case. We calculate TER as our experiments with greedy decoding for CTC AM showed good results with the visual adaptation, and the implicit LM of S2S was also quite strong, as we show below. The 'dev' set is a tougher set than the 'test' set in the How-To dataset.

	A CTC	A+V CTC	A S2S	A+V S2S
TER dev	15.2	14.1	18.4	16.8
TER test	13.6	13.1	16.3	15.7
PPL* dev	113.6	80.6	1.38	1.37
PPL* test	112.0	72.0	1.05	1.05

Table 1. Results for Audio(A) and Audio-Visual(A+V) adaptation with the How-To data

5.2. CTC vs. S2S

We compare the audio-only CTC and S2S models trained on WSJ and 90 h subset of How-To in Table 2. TER with CTC and S2S on WSJ is a strong baseline compared to prior work [23, 34, 25, 30]. We see a huge disparity in ASR for clean prepared data (WSJ) and real application data (How-To) as discussed in Section 3. We also note that for noisy data like How-To, CTC required 5 times as much data to make the error rates between S2S and CTC comparable.

	CTC	S2S
WSJ	6.9	7.9
How-To	18.5	15.3

Table 2. TER of CTC and S2S on WSJ (rval92), How-To(test set)

Table 3 shows an example from the How-To test set which highlights the fundamental difference between the CTC and S2S model when applied to "real" data: while the TER for both models is similar in value, it stems from different types of mistakes. The CTC model's output is closer to the acoustics and produces characters even for corrections and false starts (even though those were also not present in the training transcripts, according to our analysis); there is a need for language model re-scoring. The S2S model however takes more liberty in character prediction and the resulting output is a combination of the acoustics, and the language model, which has been trained to ignore false starts. Even on an S2S system that has been trained on 90 h only, it appears that the inbuilt language model is strong and does not produce out of vocabulary words like CTC does, similar to the effect we observed when decoding CTC AMs with an RNN LM on the Switchboard corpus [35].

Spoken	now it does only say for do- or doesn't even say for	
	dogs or cats it's neither	
Reference	now doesn't even say dogs or cats it says neither	
Greedy	now it does only saye for dog or toat use a dogs or	
CTC	cats s niter	
S2S	now it doesn't we say for a dog or that use a dogs	
	or cats so is night or	
AV S2S	now it doesn't leave safer dog or it does use a dogs	
	or cat so in night or	

Table 3. Typical transcription on How-To test set: the CTC model is close to acoustics even during a correction (following "do-"), while S2S keeps to the style of the reference, which is itself an abstraction of the spoken content. Currently, there is little semantic difference between regular and adapted (AV) S2S (or CTC) output.

5.3. WSJ vs. How-To in S2S

Figure 3 compares reference length to hypothesis length for 40 short and long WSJ and How-To test utterances. On WSJ, the range of lengths of short and long utterances are similar, and reference and hypothesis follow each other closely. On How-To, hypothesis prediction is very unstable and the model makes a lot of mistakes, even breaks completely at times. Length of the hypothesis is greater than the length of reference for short utterances, while it is lesser for longer utterances. As seen from the example in Table 3, the output of the S2S model is much closer to the reference transcript. The model learns a form of length normalization over the entire dataset hence performs badly on short and long utterances. We use the length normalization factor during decoding to stabilize the output of the S2S model and get absolute improvements of 2% (dev) and 1% (test) for the non-adapted case, which is slightly better than the adapted case and shows that adaptation stabilizes model performance. With our model we see no need for normalization on WSJ, different from [32].



Fig. 3. Length normalization by S2S for WSJ and How-To

6. CONCLUSIONS AND FUTURE WORK

In this paper, we describe and compare two end-to-end models that learn semantically relevant context information from the visual channel of a video, and use it to improve speech transcription. We show how to adapt a CTC bidirectional LSTM acoustic model and a S2S model to the visual semantic features. We compare the behavior of CTC and S2S models on a clean (WSJ) and noisy (How-To) dataset, and see that CTC output tends to be very close to the acoustics of an utterance, while S2S output appears to be closer to the style of the transcriptions. Similar to [36], we find that S2S approaches are maybe surprisingly robust against "real-world" data, that has not been carefully prepared for speech recognition experiments. On the WSJ dataset, our system outperforms previous S2S implementations like [30]. Next, we would try more adaptation strategies like the MLP shift with S2S model and re-scoring output of CTC with a char-RNN LM to form an all neural model.

7. ACKNOWLEDGMENTS

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. Specifically, it used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

8. REFERENCES

- Y. Miao and F. Metze, "Open-domain audio-visual speech recognition: A deep learning approach.," in *Proc. INTER-SPEECH*, 2016, pp. 3414–3418.
- [2] A. Gupta, Y. Miao, L. Neves, and F. Metze, "Visual features for context-aware speech recognition," in *Proc. ICASSP*. IEEE, 2017, pp. 5020–5024.
- [3] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *Proc. CVPR*, 2015, pp. 3156–3164.
- [4] H. Liao, E. McDermott, and A. Senior, "Large scale deep neural network acoustic modeling with semi-supervised training data for youtube video transcription," in *Proc. ASRU*. IEEE, 2013, pp. 368–373.
- [5] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *Proc. ASRU*. 2013, pp. 55–59, IEEE.
- [6] V. Gupta, P. Kenny, P. Ouellet, and T. Stafylakis, "I-vectorbased speaker adaptation of deep neural networks for french broadcast audio transcription," in *Proc. ICASSP*. IEEE, 2014.
- [7] A. Senior and I. Lopez-Moreno, "Improving dnn speaker independence with i-vector inputs," in *Proc. ICASSP*. IEEE, 2014.
- [8] Y. Miao, H. Zhang, and F. Metze, "Speaker adaptive training of deep neural network acoustic models using i-vectors," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 23, no. 11, pp. 1938–1949, Nov. 2015.
- [9] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," *IEEE transactions on acoustics, speech, and signal processing*, vol. 37, no. 3, pp. 328–339, 1989.
- [10] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proc. INTERSPEECH*. ISCA, 2015.
- [11] K. Vesely, S. Watanabe, M. Karafi, J. Honza, et al., "Sequence summarizing neural network for speaker adaptation," in *Proc. ICASSP*. IEEE, 2016, pp. 5315–5319.
- [12] Y. Miao and F. Metze, "Distance-aware DNNs for robust speech recognition," in *Proc. INTERSPEECH*, Dresden, Germany, Sept. 2015, ISCA.
- [13] S. Dupont and J. Luettin, "Audio-visual speech modeling for continuous speech recognition," *IEEE Transactions on Multimedia*, vol. 2, no. 3, pp. 141–151, 2000.
- [14] C. Bregler, H. Hild, S. Manke, and A. Waibel, "Improving connected letter recognition by lipreading," in *Proc. ICASSP*. IEEE, 1993, vol. 1, pp. 557–560.
- [15] G. Potamianos, J. Luettin, and C. Neti, "Hierarchical discriminant features for audio-visual lvcsr," in *Proc. ICASSP.* IEEE, 2001, vol. 1, pp. 165–168.
- [16] B. P. Yuhas, M. H. Goldstein, and T. J. Sejnowski, "Integration of acoustic and visual speech signals using neural networks," *IEEE Communications Mag.*, vol. 27, no. 11, pp. 65–71, 1989.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1097–1105.

- [18] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Proc. NIPS*, 2014, pp. 487–495.
- [19] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, pp. 369–376.
- [20] H. Sak, A. Senior, K. Rao, O. Irsoy, A. Graves, F. Beaufays, and J. Schalkwyk, "Learning acoustic frame labeling for speech recognition with recurrent neural networks," in *Proc. ICASSP.* IEEE, 2015, pp. 4280–4284.
- [21] 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.
- [22] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, "A neural probabilistic language model," in *JMLR*, 2003.
- [23] A. Y. Hannun, A. L. Maas, D. Jurafsky, and A. Y. Ng, "Firstpass large vocabulary continuous speech recognition using bidirectional recurrent DNNs," arXiv preprint: 1408.2873, 2014.
- [24] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. NIPS*, 2014.
- [25] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. ICASSP.* IEEE, 2016, pp. 4945–4949.
- [26] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell," arXiv preprint arXiv:1508.01211, 2015.
- [27] Y. Miao, L. Jiang, H. Zhang, and F. Metze, "Improvements to speaker adaptive training of deep neural networks," in 2014 IEEE Spoken Language Technology Workshop. IEEE, 2014.
- [28] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-toend continuous speech recognition using attention-based recurrent NN: First results," *arXiv preprint arXiv*:1412.1602, 2014.
- [29] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," *arXiv* preprint arXiv:1508.04025, 2015.
- [30] S. Kim, T. Hori, and S. Watanabe, "Joint ctc-attention based end-to-end speech recognition using multi-task learning," in *Proc. ICASSP.* IEEE, 2017, pp. 4835–4839.
- [31] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 41–48.
- [32] J. Chorowski and N. Jaitly, "Towards better decoding and language model integration in sequence to sequence models," *arXiv preprint arXiv:1612.02695*, 2016.
- [33] G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush, "OpenNMT: Open-Source Toolkit for Neural Machine Translation," ArXiv e-prints.
- [34] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *Proc. ICML*, 2014.
- [35] T. Zenkel, R. Sanabria, F. Metze, J. Niehues, M. Sperber, S. Stüker, and A. Waibel, "Comparison of decoding strategies for ctc acoustic models," in *Proc. INTERSPEECH*, Stockholm, Sweden, Aug. 2017, ISCA, Accepted.
- [36] C. Chiu, A. Tripathi, K. Chou, C. Co, N. Jaitly, D. Jaunzeikare, A. Kannan, P. Nguyen, H. Sak, A. Sankar, J. Tansuwan, N. Wan, Y. Wu, and X. Zhang, "Speech recognition for medical conversations," *CoRR*, vol. abs/1711.07274, 2017.