

FUNDAMENTAL FREQUENCY FEATURE NORMALIZATION AND DATA AUGMENTATION FOR CHILD SPEECH RECOGNITION

Gary Yeung, Ruchao Fan, and Abeer Alwan

Dept. of Electrical and Computer Engineering, University of California, Los Angeles, USA

ABSTRACT

Automatic speech recognition (ASR) systems for young children are needed due to the importance of age-appropriate educational technology. Because of the lack of publicly available young child speech data, feature extraction strategies such as feature normalization and data augmentation must be considered to successfully train child ASR systems. This study proposes a novel technique for child ASR using both feature normalization and data augmentation methods based on the relationship between formants and fundamental frequency (f_o). Both the f_o feature normalization and data augmentation techniques are implemented as a frequency shift in the Mel domain. These techniques are evaluated on a child read speech ASR task. Child ASR systems are trained by adapting a BLSTM-based acoustic model trained on adult speech. Using both f_o normalization and data augmentation results in a relative word error rate (WER) improvement of 19.3% over the baseline when tested on the OGI Kids' Speech Corpus, and the resulting child ASR system achieves the best WER currently reported on this corpus.

Index Terms— child speech recognition, fundamental frequency, feature normalization, data augmentation

1. INTRODUCTION

The development of effective child automatic speech recognition (ASR) systems has become important in recent years. For example, the advancement of child ASR can facilitate the development of teaching and assessment tools for children in educational settings [1, 2, 3] using interactive systems such as social robots [4, 5, 6]. This is especially relevant for kindergarten-aged children who are just learning to read, write, or type and rely on speech to interact with technology. Yet, ASR systems for young children still perform quite poorly when compared to adult ASR [6, 7].

One of the major hurdles facing the development of effective child ASR systems is the lack of publicly available young child speech databases. This is especially a concern in an era where deep learning, which requires many hours of

training data, is rapidly becoming the primary method of developing ASR systems, using data intensive acoustic models such as bidirectional long-short term memory (BLSTM) networks. Hence, many child ASR systems complement young child speech with older child speech or even adult speech for training data. However, there is a large acoustic mismatch between child and adult speech, further complicated by the fact that children's speech acoustics change quite dramatically as they grow [8, 9]. These changes include formants and fundamental frequency (f_o) [8, 9], two defining features of the speech signal, especially for vowels. Furthermore, as the age difference between training and testing speakers grows, ASR performance degrades rapidly [7].

One common strategy to account for the acoustic mismatch between speakers is frequency normalization. This approach attempts to warp the frequency spectra of an utterance given a normalization factor for the utterance and a target acoustic space. For instance, vocal tract length normalization (VTLN) warps the frequency spectra using a maximum likelihood approach and can be implemented in several ways with varying degrees of success [10, 11, 12]. An alternative approach is to use acoustically relevant speech parameters as normalization factors such as subglottal resonances (SGRs) [13], the third spectral peak or formant frequency [14], and f_o [15].

Another strategy is to augment the training data by creating additional speech-like features for training data. There are a number of ways to implement this augmentation such as manipulating the frequency scaling or adding noise [16, 17]. While data augmentation has not been as readily explored for child speech compared to adult speech, some techniques that have been evaluated include adding noise and reverberation [18] and using out-of-domain adult data [19].

In our previous study, we proposed an f_o -based normalization technique for child ASR [15]. In that study, a number of ASR systems were trained using speech from children of various ages and tested using kindergarten-aged children. While that study demonstrated that f_o normalization was effective for child ASR, the experiments performed assumed that only older child speech was available to train a young child ASR system. In a more practical situation, we may expect that some amount of in-domain child speech would be available as training data. Furthermore, we expect to be able

This work was supported in part by National Science Foundation (NSF) Grant #1734380.

to generate additional training data through effective data augmentation techniques.

In this study, we propose a new data augmentation method for training child ASR systems based on the f_o normalization method we proposed in [15]. This method is capable of generating speech-like features that adhere to the physical properties of speech defined by the relationship between vowel formants and f_o . We show that this data augmentation method is capable of improving child ASR systems adapted from BLSTM acoustic models trained on adult speech. Additionally, we demonstrate that this data augmentation method can be used alongside f_o normalization for further improvement.

The remainder of the paper is organized as follows. Section 2 reviews the f_o normalization technique proposed in [15] and formulates the data augmentation technique. Section 3 describes the databases and experimental setup. Section 4 presents the experimental results. Section 5 concludes the paper with a summary and considerations for future work.

2. NORMALIZATION TECHNIQUE

2.1. Relationship Between f_o and Formants

Several past studies have revealed meaningful connections between f_o and the first several formant frequencies ($F1$, $F2$, $F3$, ...) of vowels. In Bark scale, the tonotopic distances between adjacent formants ($F(x+1) - Fx$ for $x \in \{1, 2, 3, \dots\}$), along with the tonotopic distance between the first formant and f_o ($F1 - f_o$), have proven to be effective at modeling human vowel perception and representing the vowel space [20, 21]. Furthermore, these tonotopic distances can be equivalently represented by the difference between any formant and f_o ($Fx - f_o$ for $x \in \{1, 2, 3, \dots\}$). This suggests that the vowel space can be modeled as a linear relationship (in Bark scale) between formants and f_o . Additionally, studies have found that f_o and formant locations depend on one another in both vowel production [22] and vowel perception [23]. A more detailed examination of these relationships can be found in [15].

2.2. Normalization Procedure

The frequency normalization technique we proposed for kindergarten ASR in [15] was designed to be applied to any spectral-based feature such as Mel-frequency cepstral coefficients (MFCCs). The linear relationship between f_o and vowel formants is exploited to formulate a normalization technique using median f_o as the only normalization factor. As such, the technique attempts to map a speaker’s acoustic space, governed by the speaker’s f_o , to some target space, defined by a predetermined target f_o . Notably, Mel scale was used in [15] instead of Bark scale as the two scales are highly correlated, but either scale can be used in practice. For consistency, this study will also use the Mel scale.

To perform the f_o -based normalization, a default f_o , denoted as $f_{o,def}$, must first be chosen. Before feature computation, f_o extraction is performed across the entire utterance using a reliable f_o detection algorithm. For this study, we will use the multi-band summary correlogram (MBSC) pitch detection algorithm [24].

The median f_o across all voiced frames of the utterance, denoted as $f_{o,utt}$, is then chosen as the normalization factor for the utterance. The discrete Fourier transform (DFT) of the feature extraction procedure is warped as follows:

$$f_{norm} = f_{orig} - (f_{o,utt} - f_{o,def}) \quad (1)$$

where all frequencies are in Mel scale, f_{norm} is a normalized frequency corresponding to some DFT index, and f_{orig} is the frequency from the original speech spectrum mapped to f_{norm} .

An example of this f_o normalization is shown in Figure 1. The Mel filter bank outputs of an 18 year old male and a 7 year old male saying the vowel /i/ are displayed both with and without f_o normalization. When both utterances are normalized to $f_{o,def} = 100$ Hz, the Mel filter bank outputs become more similar.

2.3. Data Augmentation Procedure

While the f_o normalization procedure attempted to reduce variability between speakers by using Eq. 1, fixing $f_{o,def}$ to a default value, and adjusting $f_{o,utt}$, an alternative procedure can use Eq. 1 to create variability rather than reduce it. This can be accomplished by extracting features multiple times from the same speech utterance while adjusting $f_{o,def}$. By perturbing the feature extraction, we can augment the training data by generating acoustic features that are consistent with the structure of speech defined by the tonotopic distances. This can be especially useful for training deep neural networks that require a large amount of speech data. We will refer to this technique as f_o perturbation. When combined with the f_o normalization procedure, we can simultaneously remove larger inter-speaker variabilities while perturbing the features for additional training data.

3. DATABASE AND EXPERIMENTAL SETUP

3.1. Database

Two child speech databases were used in this study. The first one was the OGI Kids’ Speech Corpus [25]. This corpus contains approximately 100 speakers per American educational grade level, from kindergarten to 10th grade. Utterances were recorded with a sampling rate of 16 kHz (8 kHz bandwidth). Both scripted and spontaneous styles of speech were recorded from each speaker. In this study, we used the sentence utterances from the scripted speech recordings, which consisted of a total of 10,072 sentence utterances from children in grades




Fig. 1: Mel filter bank outputs of an 18 year old male (solid) and 7 year old male (dashed) saying the vowel /i/, computed with 15 filters and a frequency range of 20 Hz to 6 kHz. The filter outputs are computed both without (left) and with (right) f_o normalization. When normalization is applied, default f_o is chosen to be $f_{o,def} = 100$ Hz. The 18 year old male had $f_{o,utt} = 106$ Hz, and the 7 year old male had $f_{o,utt} = 270$ Hz. The Mel filter bank outputs computed with f_o normalization are much more well-aligned than the filter outputs computed without f_o normalization.

kindergarten through 5th grade. Approximately 70% of the utterances in each grade were used as training data for a total of 7,051 training utterances. The remaining utterances were used for testing.

The second child speech database was the CMU Kids Corpus [26]. This corpus contains 76 speakers between 1st and 3rd grade with two additional speakers from kindergarten or 6th grade for a total of 78 speakers. Utterances were recorded with a sampling rate of 16 kHz. A total of 5,180 read sentence utterances were recorded across these speakers. Exactly 70% of the utterances in this corpus were used as training data for a total of 3,626 training utterances. The remaining utterances were used for testing.

For adult data, the LibriSpeech ASR Corpus was used [27]. For our experiments, we used all the training data from the corpus, which contains 960 hours of adults reading audio books in clean and noisy conditions, for adult model training.

3.2. Feature Extraction

The baseline features used for the ASR experiments were 13-dimensional MFCCs with a window size of 25 ms and shift of 10 ms. These MFCCs were extracted using a 512-point DFT, 23 Mel filters, and bandwidth from 20 Hz to 8 kHz.

The second set of features was similar to the baseline features except the DFT was normalized by the procedure presented in Section 2.2. The default f_o was chosen to be $f_{o,def} = 100$ Hz representing an adult male f_o , and $f_{o,utt}$ was chosen to be median f_o across the utterance, estimated using MBSC pitch detection. As the value of $f_{o,utt}$ reached as high as 300 Hz, the bandwidth for these features was limited to 6.2

kHz to compensate for the frequency shift upwards. Notably, this limits the maximum frequency shift of the f_o normalization procedure to approximately 250 Mels. Both baseline and f_o normalized features were extracted from all utterances.

When applying f_o perturbation, $f_{o,def}$ was adjusted such that the DFT was shifted by ± 20 , ± 40 , and ± 60 Mels. Along with the original features, this multiplies the amount of training data by 7. Data augmentation was applied to the CMU Kids and OGI Kids’ training datasets.

Implementation of the feature extraction is relatively simple when using Eq. 1 to compute the DFT frequency shift. We let $f_{o,def} \in \{58.52, 72.10, 85.93, 100.00, 114.32, 128.90, 143.74\}$ Hz where every value in the set is used if performing data augmentation and only 100 Hz is used if performing a standard feature extraction. Similarly, we let $f_{o,utt} = 100$ Hz when performing a non-normalized feature extraction and choose $f_{o,utt}$ as the median f_o across the utterance when normalizing.

3.3. Experimental Setup

For both the baseline and f_o normalized features, an adult ASR system was first trained using the LibriSpeech training set. The acoustic model was a 3-layer BLSTM network with 512 cells in each direction followed by a feed-forward layer that mapped the output of the BLSTMs to senone probabilities. The input to the BLSTM was 7 frames (3 frames forward and backward in time) to form a 91-dimensional input for the acoustic model. The output was approximately 5,700 senone probabilities. The acoustic model was based on PyKaldi2 [28], while decoding used the Kaldi Speech Recog-

Table 1: Word error rates (WERs) of the child ASR experiment using a BLSTM-based acoustic model adapted from adult speech. The left two columns indicate whether f_o normalization (“Norm?”) and data augmentation using f_o perturbation (“Aug?”) were used. WERs for both CMU Kids and OGI Kids’ are reported in the latter columns.

Norm?	Aug?	CMU Kids	OGI Kids’
No	No	16.88	6.84
Yes	No	16.93	6.50
No	Yes	16.63	5.85
Yes	Yes	16.47	5.52

nition Toolkit [29].

The adult ASR systems were adapted to child speech using either the CMU Kids or OGI Kids’ training data. The adult ASR system trained with f_o normalization was adapted using child speech features that were extracted with f_o normalization. Similarly, the adult ASR system trained without normalization was adapted using features that were extracted without normalization. Additionally, data augmentation using f_o perturbation was applied on the child training datasets.

A 4-gram language model (LM) trained on Project Gutenberg books was chosen for decoding. This LM is one of the default language models included in Kaldi’s LibriSpeech recipe [29]. Adapted ASR systems were evaluated using the corresponding CMU Kids or OGI Kids’ testing datasets. The ASR system trained with f_o normalization was tested using features extracted with f_o normalization.

4. RESULTS AND DISCUSSION

The results of the child ASR experiments described in Section 3 are shown in Table 1. The top row displays the word error rate (WER) of the baseline system (i.e., no normalization or augmentation). The second row displays the WER of the system using f_o normalization. The third row displays the system adapted with f_o perturbed child speech features. Finally, the last row uses both f_o normalization and f_o perturbed adaptation data. Additionally, we performed the experiment once more using the acoustic model trained only on LibriSpeech. This system achieved a WER of 37.49% for CMU Kids and 59.70% for OGI Kids’, significantly worse than any WER in Table 1, which demonstrates the importance of child speech adaptation data.

When using adaptation and only applying f_o normalization, the performance of the OGI Kids’ system saw a small improvement over the baseline, which has no normalization or data augmentation, but not enough to be significant. Applying f_o perturbation to the training data, a more substantial improvement was achieved from 6.84% to 5.85%. However, when using the CMU Kids system, no major improvements

were observed.

Using both f_o normalization and f_o perturbation resulted in the best performing ASR system for both testing sets. The OGI Kids’ testing set saw a relative WER improvement of 19.3%, reducing the WER to 5.52%, and this result is statistically significant at $p < 0.001$. However, the CMU Kids testing set only saw a relative WER improvement of 2.4%.

While the OGI Kids’ testing set saw a larger improvement than the CMU Kids testing set, this result may be expected. We note that the CMU Kids testing set had a narrower age range (approximately 6-9 years old excluding the two outlier children) compared to the OGI Kids’ testing set (approximately 5-11 years old). The f_o normalization method has been shown to produce larger improvements when the range of ages used in training and testing data is wider [15]. A similar phenomenon may be occurring for f_o perturbation. That is, since there is less variability in the CMU Kids testing set, adding additional variability to the training set through f_o perturbation was unnecessary to train the BLSTM. Meanwhile, with the larger variability of the OGI Kids’ dataset, both f_o normalization and f_o perturbation proved helpful. Furthermore, these techniques may extend to ASR systems using the full OGI Kids’ scripted speech dataset, which has been reported to have a WER of 10.8% [18]. Preliminary experiments suggest that f_o perturbation also performs better than VTLN-based data augmentation.

5. CONCLUSION

This study proposes a new data augmentation method for training child ASR systems based on the f_o normalization method proposed in [15]. Both normalization and data augmentation methods adhere to the physical properties of speech defined by the relationship between vowel formants and f_o . The two methods can be formulated as a simple shift of DFT bin frequencies in the Mel domain and are implemented by manipulating $f_{o,def}$ and $f_{o,utt}$ in Eq. 1. Child ASR systems were trained using these methods by adapting from a BLSTM acoustic model trained on adult speech. When using both f_o normalization and f_o perturbation, a 19.3% relative improvement was observed on the OGI Kids’ Speech Corpus. However, a less substantial improvement was observed on the CMU Kids Corpus. This suggests, that both f_o normalization and f_o perturbation are more effective when the age range of the speakers is large.

There are a number of possible directions for future work. We plan to evaluate these methods on ASR systems for both children and adults for a situation with increased age and f_o variability between speakers. We also plan to apply these systems to other child ASR applications such as conversational speech, educational applications, and clinical applications.

6. REFERENCES

- [1] J. Tepperman, J. Silva, A. Kazemzadeh, H. You, S. Lee, A. Alwan, and S. Narayanan, "Pronunciation Verification of Children's Speech for Automatic Literacy Assessment," in *Proc. INTERSPEECH*, 2006, pp. 845–848.
- [2] H. T. Bunnell, D. M. Yarrington, and J. B. Polikoff, "STAR: Articulation Training for Young Children," in *Proc. ICSLP*, 2000, pp. 85–88.
- [3] G. Yeung, A. Afshan, K. E. Ozgun, K. Kaewtip, S. M. Lulich, and A. Alwan, "Predicting Clinical Evaluations of Children's Speech with Limited Data Using Exemplar Word Template References," in *Proc. SLATE*, 2017, pp. 161–166.
- [4] S. Spaulding, H. Chen, S. Ali, M. Kulinski, and C. Breazeal, "A Social Robot System for Modeling Children's Word Pronunciation," in *Proc. AAMAS*, 2018, pp. 1658–1666.
- [5] G. Yeung, A. L. Bailey, A. Afshan, M. Q. Pérez, A. Martin, S. Spaulding, H. W. Park, A. Alwan, and C. Breazeal, "Towards the Development of Personalized Learning Companion Robots for Early Speech and Language Assessment," in *Proc. AERA*, 2019.
- [6] J. Kennedy, S. Lemaignan, C. Montassier, P. Lavalade, B. Irfan, F. Papadopoulos, E. Senft, and T. Belpaeme, "Child Speech Recognition in Human-Robot Interaction: Evaluations and Recommendations," in *Proc. ACM/IEEE HRI*, 2017, pp. 82–90.
- [7] G. Yeung and A. Alwan, "On the Difficulties of Automatic Speech Recognition for Kindergarten-Aged Children," in *Proc. INTERSPEECH*, 2018, pp. 1661–1665.
- [8] S. Lee, A. Potamianos, and S. Narayanan, "Acoustics of Children's Speech: Developmental Changes of Temporal and Spectral Parameters," *J. Acoust. Soc. Am.*, vol. 105, no. 3, pp. 1455–1468, 1999.
- [9] H. K. Vorperian and R. D. Kent, "Vowel Acoustic Space Development in Children: A Synthesis of Acoustic and Anatomic Data," *J. Speech. Lang. Hear.*, vol. 50, no. 6, pp. 1510–1545, 2007.
- [10] G. Stemmer, C. Hacker, S. Steidl, and E. Nöth, "Acoustic Normalization of Children's Speech," in *Proc. EUROSPEECH*, 2003, pp. 1313–1316.
- [11] X. Cui and A. Alwan, "MLLR-Like Speaker Adaptation Based on Linearization of VTLN with MFCC Features," in *Proc. INTERSPEECH*, 2005, pp. 273–276.
- [12] R. Serizel and D. Giuliani, "Vocal Tract Length Normalization Approaches to DNN-Based Children's and Adults' Speech Recognition," in *Proc. SLT*, 2014, pp. 135–140.
- [13] J. Guo, R. Paturi, G. Yeung, S. M. Lulich, H. Arsikere, and A. Alwan, "Age-Dependent Height Estimation and Speaker Normalization for Children's Speech Using the First Three Subglottal Resonances," in *Proc. INTERSPEECH*, 2015, pp. 1665–1669.
- [14] X. Cui and A. Alwan, "Adaptation of Children's Speech with Limited Data Based on Formant-Like Peak Alignment," *Comput. Speech Lang.*, vol. 20, no. 4, pp. 400–419, 2006.
- [15] G. Yeung and A. Alwan, "A Frequency Normalization Technique for Kindergarten Speech Recognition Inspired by the Role of fo in Vowel Perception," in *Proc. INTERSPEECH*, 2019, pp. 6–10.
- [16] X. Cui, V. Goel, and B. Kingsbury, "Data augmentation for deep neural network acoustic modeling," in *Proc. ICASSP*, 2014, pp. 5582–5586.
- [17] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. INTERSPEECH*, 2019, pp. 2613–2617.
- [18] F. Wu, L. P. Garcia, D. Povey, and S. Khudanpur, "Advances in automatic speech recognition for child speech using factored time delay neural network," in *Proc. INTERSPEECH*, 2019, pp. 1–5.
- [19] J. Fainberg, P. Bell, M. Lincoln, and S. Renals, "Improving children's speech recognition through out-of-domain data augmentation," in *Proc. of INTERSPEECH*, 2016, pp. 1598–1602.
- [20] A. K. Syrdal and H. S. Gopal, "A Perceptual Model of Vowel Recognition Based on the Auditory Representation of American English Vowels," *J. Acoust. Soc. Am.*, vol. 79, no. 4, pp. 1086–1100, 1986.
- [21] H. Traunmüller, "Perceptual Dimension of Openness in Vowels," *J. Acoust. Soc. Am.*, vol. 69, no. 5, pp. 1465–1475, 1981.
- [22] S. Barreda and T. M. Nearey, "The Perception of Formant-Frequency Range is Affected by Veridical and Judged Fundamental Frequency," in *Proc. Meetings on Acoustics*, 2013, vol. 19, p. 060197.
- [23] S. Barreda and T. M. Nearey, "The Direct and Indirect Roles of Fundamental Frequency in Vowel Perception," *J. Acoust. Soc. Am.*, vol. 131, no. 1, pp. 466–477, 2012.
- [24] L. N. Tan and A. Alwan, "Multi-Band Summary Correlogram-Based Pitch Detection for Noisy Speech," *Speech Comm.*, vol. 55, no. 7-8, pp. 841–856, 2013.
- [25] K. Shobaki, J.-P. Hosom, and R. A. Cole, "The OGI Kids' Speech Corpus and Recognizers," in *Proc. ICSLP*, 2000, pp. 258–261.
- [26] M. Eskenazi, J. Mostow, and D. Graff, "The CMU Kids Speech Corpus LDC97S63," 1997.
- [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. IEEE ICASSP*, 2015, pp. 5206–5210.
- [28] L. Lu, X. Xiao, Z. Chen, and Y. Gong, "PyKaldi2: Yet another speech toolkit based on Kaldi and PyTorch," *arXiv:1907.05955*, 2019.
- [29] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlíček, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Veselý, "The Kaldi Speech Recognition Toolkit," in *Proc. IEEE ASRU*, 2011.