



Voice Conversion Using Sequence-to-Sequence Learning of Context Posterior Probabilities

Hiroyuki Miyoshi, Yuki Saito, Shinnosuke Takamichi, and Hiroshi Saruwatari

Graduate School of Information Science and Technology, The University of Tokyo,
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

mathmal306@gmail.com

{yuuki.saito, shinnosuke.takamichi, hiroshi.saruwatari}@ipc.i.u-tokyo.ac.jp

Abstract

Voice conversion (VC) using sequence-to-sequence learning of context posterior probabilities is proposed. Conventional VC using shared context posterior probabilities predicts target speech parameters from the context posterior probabilities estimated from the source speech parameters. Although conventional VC can be built from non-parallel data, it is difficult to convert speaker individuality such as phonetic property and speaking rate contained in the posterior probabilities because the source posterior probabilities are directly used for predicting target speech parameters. In this work, we assume that the training data partly include parallel speech data and propose sequence-to-sequence learning between the source and target posterior probabilities. The conversion models perform non-linear and variable-length transformation from the source probability sequence to the target one. Further, we propose a joint training algorithm for the modules. In contrast to conventional VC, which separately trains the speech recognition that estimates posterior probabilities and the speech synthesis that predicts target speech parameters, our proposed method jointly trains these modules along with the proposed probability conversion modules. Experimental results demonstrate that our approach outperforms the conventional VC.

Index Terms: voice conversion, context posterior probabilities, sequence-to-sequence learning

1. Introduction

Voice conversion (VC) is a technique for converting para- and non-linguistic information while keeping linguistic information. VC is used in a variety of applications such as speech enhancement [1, 2] and language education for non-native speakers [3]. It is mainly classified into two types: text-independent VC and text-dependent VC. Text-independent VC directly predicts target speech parameters from the source speech parameters, and acoustic models such as Gaussian mixture models [4, 5] or deep neural network [6] are trained using only speech data. Since the models are often trained using parallel speech data, the conversion quality of the performance is typically highly accurate. However, in most cases, parallel data is not readily available. Text-dependent VC [7, 8], in contrast, converts speech parameters through textual information. This type consists of two modules: speech recognition that estimates the textual information from the source speech, and speech synthesis that predicts target speech from the textual information. Basically, parallel data are not required to build the VC, and the training data are easily available. However, the conversion units of this method are rougher (e.g., phoneme, word, or other linguistic units) than those of text-independent VC (e.g., frame). VC using shared context posterior probabilities [9] is classified

in text-dependent VC, but the conversion unit is frame level. The context posterior probability of source speech parameters is estimated frame by frame, and then the target speech parameters are predicted from the estimated posterior probabilities. This VC is interpreted as *soft* text-dependent VC and can be extended to cross-lingual text-to-speech [10, 11]. However, it cannot convert speaker individuality (e.g., speaking rates and phonetic properties) included in the context posterior probabilities because the posterior probabilities of the source speech are directly used for predicting target speech parameters.

In light of the above, we propose a sequence-to-sequence learning of the context posterior probabilities. Assuming that the training data partly include parallel speech data (parallel utterances of phrases), we build an encoder-decoder model [12] that converts the posterior probabilities of the source speech parameters into those of the target speech parameters. The proposed posterior probability conversion module is inserted between conventional speech recognition and synthesis. When we do not build the conversion model or do not have parallel data, conventional VC [9] is available. Further, we propose a joint training algorithm. Whereas the conventional VC [9] separately trains speech recognition and speech synthesis, our approach jointly trains these modules (like auto-encoding) and the proposed conversion module. We found through experiment that the proposed methods outperform the conventional VC [9].

2. VC Using Shared Context Posterior Probabilities

Conventional VC using shared context posterior probabilities [9] contains two modules: speech recognition and speech synthesis. They are separately trained, and the voice conversion is performed by concatenating them. Figure 1 shows an example of context posterior probabilities. The upper and middle parts of Fig. 2 show the details of these processes.

2.1. Training Stage

Recognition models estimate context posterior probability sequence from the speech parameter sequence. Let $\mathbf{x} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_{T_x}^\top]^\top$ and $\mathbf{y} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_{T_y}^\top]^\top$ be source and target speech parameter sequences, respectively. \mathbf{x}_t and \mathbf{y}_t are the parameters at frame t . T_x and T_y are their frame lengths. Also, let $\mathbf{l}^{(x)} = [\mathbf{l}_1^{(x)}, \dots, \mathbf{l}_{T_x}^{(x)}]^\top$ and $\mathbf{l}^{(y)} = [\mathbf{l}_1^{(y)}, \dots, \mathbf{l}_{T_y}^{(y)}]^\top$ be the context label sequence (such as quin-phone) corresponding to \mathbf{x} and \mathbf{y} , respectively. Speaker-independent neural network $\mathbf{R}(\cdot)$ is trained using speech data including \mathbf{x} and \mathbf{y} , and the training criterion is minimizing the cross entropy $L_C(\mathbf{l}_x, \mathbf{R}(\mathbf{x}))$.

Synthesis models predict target speech parameter sequence \mathbf{y} from the corresponding context posterior probability se-




Figure 1: An example of source (left) and target (right) context posterior probabilities.

quence $\hat{\mathbf{p}}_y$, using the trained recognition models, i.e., $\hat{\mathbf{p}}_y = \mathbf{R}(\mathbf{y})$. The target-speaker-dependent neural networks $\mathbf{G}(\cdot)$ are trained to minimize the mean squared error $L_G(\mathbf{y}, \mathbf{G}(\hat{\mathbf{p}}_y))$ between \mathbf{y} and $\mathbf{G}(\hat{\mathbf{p}}_y)$.

2.2. Conversion Stage

In conversion, the converted speech parameter sequence $\hat{\mathbf{y}}$ is predicted by concatenating speech recognition and speech synthesis, i.e., $\hat{\mathbf{y}} = \mathbf{G}(\hat{\mathbf{p}}_x) = \mathbf{G}(\mathbf{R}(\mathbf{x}))$, where $\hat{\mathbf{p}}_x$ is the context posterior probability sequence of \mathbf{x} . Note that the frame lengths of \mathbf{x} , $\hat{\mathbf{p}}_x$ and $\hat{\mathbf{y}}$ are the same, i.e., T_x .

2.3. Problems

Since the posterior probabilities estimated in speech recognition are directly used for speech synthesis, it is difficult to convert speaker individuality included in the posterior probabilities, such as the speaking rate (frame length) and phonetic characteristics (see Fig. 1). Also, improving recognition accuracy does not always improve speech quality in converted speech (except zero error in recognition).

3. Proposed VC using Sequence-to-Sequence Learning of Context Posterior Probabilities

To overcome the limitation of the conventional method, we propose an approach for converting source context posterior probabilities to target context posterior probabilities using sequence-to-sequence learning.

3.1. Sequence-to-Sequence Learning

Sequence-to-sequence learning using recurrent neural networks (RNNs) [13] can be applied to the problem that the source and target sequences have different lengths. An encoder-decoder model we adopt here maps a variable-length source sequence to a fixed-length vector, and maps the vector to the variable-length target sequence. At each frame, the source side RNN (encoder) and target side RNN (decoder) predict the source and target features of the next frame, respectively. As discussed below, we adopt this in order to convert a source posterior probability sequence to a target that has a different length. The proposed procedure is shown in the lower part of Fig. 2.

3.2. Training Stage

We propose two algorithms to perform the posterior probability conversion. The first algorithm *separately* trains speech recog-

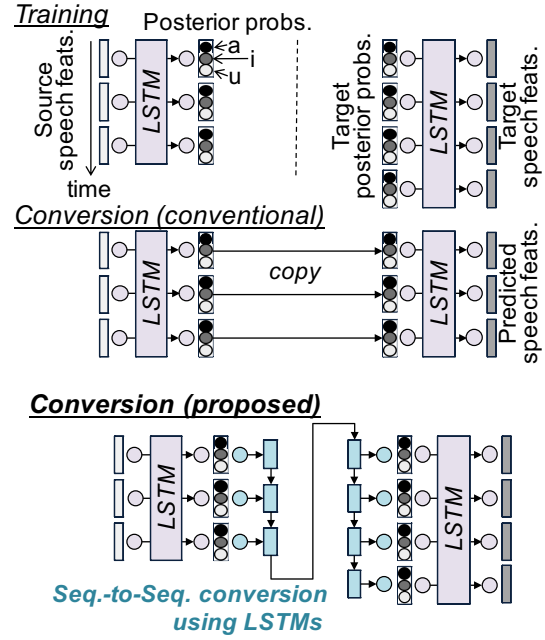


Figure 2: Training and conversion procedures of conventional and proposed VC. In the proposed VC, the source context posterior probabilities are transformed into the target posterior probabilities.

nition and synthesis the same as the conventional algorithms and *separately* trains probability conversion models using the source and target posterior probabilities. The second algorithm *jointly* trains the recognition and synthesis and trains the conversion models considering not only posterior probabilities conversion but also speech synthesis. In conversion, we concatenate the three models for converting input speech parameters.

3.2.1. Training of probability conversion models

Given the parallel sequences of source and target context posterior probabilities, we train encoder-decoder models $\mathbf{C}(\cdot)$ that convert the source and target sequences. The loss function to be minimized is as follows:

$$L(\mathbf{l}_y, \hat{\mathbf{p}}_x, \hat{\mathbf{p}}_y) = L_G(\hat{\mathbf{p}}_y, \mathbf{C}(\hat{\mathbf{p}}_x)) + L_C(\mathbf{l}_y, \mathbf{C}(\hat{\mathbf{p}}_x)). \quad (1)$$

The first term minimizes the conversion error between the predicted and target sequences. The second term minimizes the cross-entropy using \mathbf{l}_y , which was obtained by the training of $\mathbf{R}(\cdot)$, and can decrease the recognition error included in $\hat{\mathbf{p}}_y$. Our preliminary evaluation demonstrated that using this formulation results in better conversion accuracy than using only the first term.

Sequence-to-sequence learning suffers from long-term dependencies, i.e., error accumulation, so in our approach, we implement phoneme-by-phoneme probability conversion. Given the phoneme boundary of the source and target probability sequence, phoneme-independent encoder-decoder models are trained to convert the probability sequence within the current phoneme.

3.2.2. Jointly training of recognition, synthesis, and conversion

Since the final goal of the method is to minimize synthesis error, its pre-processes (i.e., recognition and conversion) must be trained by considering this error. We train speech recognition $\mathbf{R}(\cdot)$ to minimize not only recognition error but also synthesis

error (e.g., reconstruction error of auto-encoders): the loss function is $L_C(\mathbf{l}_x, \mathbf{R}(\mathbf{x})) + L_G(\mathbf{x}, \mathbf{G}(\mathbf{R}(\mathbf{x})))$. The speech synthesis $\mathbf{G}(\cdot)$ is trained in the conventional manner. We further train the conversion models to minimize not only conversion error but also synthesis error: the loss function is the sum of Eq. (1) and $L_G(\mathbf{G}(\mathbf{y}, \mathbf{C}(\hat{\mathbf{p}}(\mathbf{x}))))$.

3.3. Discussion

Text-dependent VC forcibly aligns the input speech feature segment into a single context (e.g., phoneme, syllable, or word unit) and generates output speech features from the context sequence. Although this method can flexibly transform the context sequence (e.g., variable-length conversion), it cannot avoid the effect of time quantization with mapping from speech feature segments. Meanwhile, text-independent VC with dynamic time warping (DTW) [4] aligns speech features in a frame level, but it limits the transformation of (implicitly considered) context sequence, e.g., the sequence length is fixed. The conventional method [9, 10] also corresponds to the latter because the source posterior probability is directly used for synthesizing the target speech. In comparison with these methods, since the proposed algorithm performs frame-level conversion without forced alignments, it can avoid the effect of time quantization and convert context sequence flexibly.

Joint training of recognition and synthesis, which is proposed in Section 3.2.2, is similar to auto-encoding processes [14] with the referred class labels and dual learning [15]. Therefore, we expect that these processes can be extended to the supervised learning of variational auto-encoders [16] that have not only class labels (e.g., context labels) but also the hidden variables [17, 18].

4. Experiments

4.1. Experimental Setup

Although the conventional VC [9] and proposed VC accept non-parallel speech data and partly included parallel data, we used fully parallel data in this evaluation. We prepared speech data of eight speakers taken from the ATR Japanese speech database [19]. The speaker uttered 503 phonetically balanced sentences. We built the speaker-independent speech recognition module with speaker codes by using the speech data of eight speakers including source and target speakers. We built conversion and synthesis modules by using speech data of source female and target male speakers. We used 450 sentences (subsets A to I) for the training and 53 sentences (subset J) for the evaluation. Speech signals were sampled at a rate of 16 kHz, and the shift length was set to 5 ms. The 0th-through-24th mel-cepstral coefficients were used as a spectral parameter and F_0 and 5 band-aperiodicity [20, 21] were used as excitation parameters. We used the STRAIGHT analysis-synthesis system [22] for the parameter extraction and waveform synthesis. To improve training accuracy, speech parameter trajectory smoothing [23] with a 50 Hz cutoff modulation frequency was applied to the spectral parameters in the training data. In the training phase, spectral features were normalized to have zero-mean unit-variance, and 80 % of the silent frames were removed from the training data. We used AdaGrad [24] as the optimization algorithm, setting the learning rate to 0.01. Both in speech recognition and generation stage, we used bi-directional long short-term memory (LSTM) [9, 10] and each hidden layer contained 256 units. For probability conversion, we used bi-directional LSTM in encoder, and LSTM in decoder. Each hidden layers contained 256




Figure 3: Mel-cepstral distortion of conventional VC and proposed VC integrating posterior probability conversion.

units.

Quin-phoneme was used as the context labels. For training the recognition models, we divided the quin-phoneme into five groups: previous phoneme, current phoneme, and next phoneme, and so on. The cross-entropy loss was calculated for each group, and the loss function for training the recognition models was the sum of each loss [25]. Only spectral and their delta features were used for recognition and synthesis. In the proposed method, F_0 was linearly transformed [4] first, and we modified its length using DTW between the source context posterior probability sequence and the converted posterior probability sequence. Only DTW was used for band-aperiodicity conversion. This evaluation uses reference phoneme duration for converting posterior probabilities in order to address the sequence length determination problem to which sequence-to-sequence learning is prone [26]. Given the phoneme duration of the target natural speech parameter sequence in conversion data, we performed phoneme-level probability conversion. Given the phoneme duration of the source and target speech parameters in the training and conversion data, the conversion models converted the probabilities within the current phoneme. The finally generated posterior probability sequence was then calculated by concatenating the converted probabilities.

Two evaluations were performed to compare the conventional and proposed VC. First, we evaluated the effectiveness of the proposed posterior probability conversion, and then, we evaluated the effect of the proposed joint training algorithms.

4.2. Evaluations

We discuss the effectiveness of the proposed posterior probability conversion. The separately trained modules were used here.

4.2.1. Objective Evaluation

We calculated mel-cepstral distortion between the target and converted speech parameters of the conventional VC [9] and proposed VC. DTW was used to align the target and converted parameters by the conventional VC. The difference between the two methods is the time warping method, i.e., DTW or sequence-to-sequence learning. The results (Fig. 3) clearly show that the proposed VC outperforms the conventional VC, we demonstrate that spectral distortion caused by DTW can be alleviated by the use of sequence-to-sequence learning.

4.2.2. Subjective Evaluation

To subjectively evaluate the conventional and proposed VC, we performed a preference AB test to evaluate the converted speech quality. We presented every pair of converted speech of the two sets in random order and had listeners select the speech sample that sounded better. Similarly, we performed an XAB test on the speaker individuality using the natural speech as a reference “X.” Seven listeners participated in each assessment.

The results are shown in Fig. 4. Although the proposed VC performed better in speaker similarity (Fig. 4(a)) thanks to posterior probability conversion, it degrades speech quality (Fig. 4(b)). It seems that the probability conversion caused con-




Figure 4: Results of subjective evaluation for comparing conventional VC and proposed VC integrating posterior probability conversion. Error bar indicates the 95 % confidence intervals.




Figure 5: Mel-cepstral distortion in auto-encoding case. The conventional separately trained and proposed jointly trained recognition and synthesis modules are compared.

version error that missed the phonetic properties of the source speech parameters, which is probably what resulted in the degraded quality.

4.3. Evaluation of Joint Training

4.3.1. Joint Training of Recognition and Synthesis

We evaluated the effectiveness of joint training of recognition and synthesis modules, in comparison to conventional separately trained modules [9]. We calculated mel-cepstral distortion in the auto-encoding case, i.e., reconstruction error of source speech parameters through recognition and synthesis. As shown in Fig. 5, the proposed joint training achieved better distortion than conventional separated training.

We also performed an AB test on speech quality and XAB test on speaker similarity in the VC case, as similar as in Section 4.2.2. As shown in Fig. 6, the proposed joint-training overcomes the conventional separated training in both speaker similarity and speech quality.




Figure 6: Results of objective evaluation for comparing conventional separately-trained and proposed jointly-trained recognition and synthesis modules. Error bars indicate the 95 % confidence intervals.




Figure 7: Mel-cepstral distortion of three methods: (1) conventional VC, (2) proposed VC using separately trained recognition/synthesis, and (3) proposed VC using jointly trained recognition/synthesis/conversion.




Figure 8: An example of posterior probability sequence.

4.3.2. Joint Training of Recognition, Conversion, and Synthesis

To evaluate of the joint training of recognition, conversion, and synthesis, we calculated the mel-cepstral distortion of three systems: (1) conventional VC [9], (2) separately trained modules (equal to "Proposed" in Section 4.2.1), and (3) jointly trained modules.

The results are shown in Fig. 7. The joint training scored better than conventional VC but worse than separated training. To clarify this, we show in Fig. 8 an example of probability sequence estimated through speech recognition. The separately trained recognition module outputs harder probabilities, i.e., the values are close to 0 or 1 for all frames. However, we can see that the joint training of recognition and synthesis tends to make the values soft. This requires deeper investigation, but we suspect this tendency is one of the reasons.

5. Conclusion

In this paper, we proposed voice conversion (VC) using sequence-to-sequence learning of context posterior probabilities. Since conventional VC directly uses the posterior probabilities of source speech for predicting target speech, it is difficult to convert the speaker individuality included in the posterior probabilities. To address this, we built sequence-to-sequence conversion models that convert the source context posterior probability sequence into a target one. Further, we proposed joint training algorithms for speech recognition, speech synthesis, and posterior probability conversion. Experimental results demonstrated that (1) the proposed algorithms outperformed the conventional VC in speaker similarity, and (2) joint training of recognition and synthesis outperformed the conventional VC in both speaker similarity and speech quality. As future work, we will investigate how to determine the frame length of the converted posterior probability sequence.

6. Acknowledgements

Part of this work was supported by ImPACT Program of Council for Science, Technology and Innovation (Cabinet Office, Government of Japan), SECOM Science and Technology Foundation, and JSPS KAKENHI Grant Number 16H06681.

7. References

- [1] A. B. Kain, J.-P. Hosom, X. Niu, J. P. H. van Santen, M. F.-Oken, and J. Staehely, "Improving the intelligibility of dysarthric speech," *Speech Communication*, vol. 49, no. 9, pp. 743–759, 2007.
- [2] F. Rudzicz, "Acoustic transformations to improve the intelligibility of dysarthric speech," in *Proc. SLPAT*, Edinburgh, Scotland, Jul. 2011, pp. 11–21.
- [3] S. Aryal and R. G.-Osuna, "Can voice conversion be used to reduce non-native accents?" in *Proc. ICASSP*, Florence, Italy, May 2014, pp. 7929–7933.
- [4] T. Toda, A. W. Black, and K. Tokuda, "Voice conversion based on maximum likelihood estimation of spectral parameter trajectory," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2222–2235, 2007.
- [5] Y. Stylianou, O. Cappé, and E. Moulines, "Continuous probabilistic transform for voice conversion," *IEEE Transactions on Speech and Audio Processing*, vol. 6, no. 2, pp. 131–142, 1998.
- [6] S. Desai, E. V. Raghavendra, B. Yegnanarayana, A. W. Black, and K. Prahallad, "Voice conversion using artificial neural networks," in *Proc. ICASSP*, Taipei, Taiwan, Apr. 2009, pp. 3893–3896.
- [7] A. Kain and M. W. Macon, "Spectral voice conversion for text-to-speech synthesis," in *Proc. ICASSP*, Seattle, U.S.A., May 1998, pp. 285–288.
- [8] D. Sunderman, H. Hoge, A. Bonafonte, H. Ney, A. W. Black, and S. Narayanan, "Text-independent voice conversion based on unit selection," in *Proc. ICASSP*, Toulouse, France, May 2006.
- [9] L. Sun, K. Li, H. Wang, S. Kang, and H. Meng, "Phonetic posteriorgrams for many-to-one voice conversion without parallel data training," in *Proc. ICME*, Seattle, U.S.A., Jul. 2016.
- [10] L. Sun, S. Kang, K. Li, and H. Meng, "Personalized, cross-lingual TTS using phonetic posteriorgrams," in *Proc. INTERSPEECH*, San Francisco, U.S.A., Sep. 2016, pp. 322–326.
- [11] F.-L. Xie, F. K. Soong, and H. Li, "A KL divergence and DNN-based approach to voice conversion without parallel training sentences," in *Proc. INTERSPEECH*, San Francisco, U.S.A., Sep. 2016, pp. 287–291.
- [12] K. Cho, D. Bahdanau, F. Vougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoderdecoder for statistical machine translation," in *Proc. EMNLP*, Doha, Qatar, Oct. 2014, pp. 1724–1734.
- [13] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. NIPS*, Montreal, Canada, Dec. 2014, pp. 3104–3112.
- [14] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, "Greedy layer-wise training of deep networks," in *Proc. NIPS*, Vancouver, Canada, Dec. 2006, pp. 153–160.
- [15] D. He, Y. Xia, T. Qin, L. Wang, N. Yu, T. Liu, and W.-Y. Ma, "Dual learning for machine translation," in *Proc. NIPS*, Barcelona, Spain, Dec. 2016, pp. 820–828.
- [16] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *Proc. ICLR*, Banff, Canada, Apr. 2014.
- [17] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, "Voice conversion from non-parallel corpora using variational auto-encoder," in *Proc. APSIPA ASC*, Jeju, Korea, Dec. 2016.
- [18] D. P. Kingma, D. J. Rezende, S. Mohamed, and M. Welling, "Semi-supervised learning with deep generative models," in *Proc. NIPS*, Montreal, Canada, Dec. 2014, pp. 3581–3589.
- [19] Y. Sagisaka, K. Takeda, M. Abe, S. Katagiri, T. Umeda, and H. Kuawahara, "A large-scale Japanese speech database," in *IC-SLP90*, Kobe, Japan, Nov. 1990, pp. 1089–1092.
- [20] H. Kawahara, J. Estill, and O. Fujimura, "Aperiodicity extraction and control using mixed mode excitation and group delay manipulation for a high quality speech analysis, modification and synthesis system STRAIGHT," in *MAVEBA 2001*, Firentze, Italy, Sep. 2001, pp. 1–6.
- [21] Y. Ohtani, T. Toda, H. Saruwatari, and K. Shikano, "Maximum likelihood voice conversion based on GMM with STRAIGHT mixed excitation," in *Proc. INTERSPEECH*, Pittsburgh, U.S.A., Sep. 2006, pp. 2266–2269.
- [22] H. Kawahara, I. Masuda-Katsuse, and A. D. Cheveigne, "Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based F0 extraction: Possible role of a repetitive structure in sounds," *Speech Communication*, vol. 27, no. 3–4, pp. 187–207, 1999.
- [23] S. Takamichi, K. Kobayashi, K. Tanaka, T. Toda, and S. Nakamura, "The NAIST text-to-speech system for the Blizzard Challenge 2015," in *Proc. Blizzard Challenge workshop*, Berlin, Germany, Sep. 2015.
- [24] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, pp. 2121–2159, Jul. 2011.
- [25] B. Huang, D. Ke, H. Zheng, B. Xu, Y. Xu, and K. Su, "Multi-task learning deep neural networks for speech feature denoising," in *Proc. INTERSPEECH*, Dresden, Germany, Sep. 2015, pp. 2464–2468.
- [26] W. Wang, S. Xu, and B. Xu, "First step towards end-to-end parametric TTS synthesis: Generating spectral parameters with neural attention," in *Proc. INTERSPEECH*, San Francisco, U.S.A., Sep. 2016, pp. 2243–2247.