



HAL
open science

Towards end-to-end F0 voice conversion based on Dual-GAN with convolutional wavelet kernels

Clément Le Moine, Nicolas Obin, Axel Roebel

► **To cite this version:**

Clément Le Moine, Nicolas Obin, Axel Roebel. Towards end-to-end F0 voice conversion based on Dual-GAN with convolutional wavelet kernels. EUSIPCO, 2021, Dublin (virtual), Ireland. hal-03239583

HAL Id: hal-03239583

<https://hal.science/hal-03239583v1>

Submitted on 27 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards end-to-end F0 voice conversion based on Dual-GAN with convolutional wavelet kernels

Clément Le Moine

IRCAM, CNRS, Sorbonne Université
STMS Lab,
Paris, France
lemoineveillon@ircam.fr

Nicolas Obin

IRCAM, CNRS, Sorbonne Université
STMS Lab,
Paris, France
nicolas.obin@ircam.fr

Axel Roebel

IRCAM, CNRS, Sorbonne Université
STMS Lab,
Paris, France
axel.roebel@ircam.fr

Abstract—This paper presents an end-to-end framework for the F0 transformation in the context of expressive voice conversion. A single neural network is proposed, in which a first module is used to learn F0 representation over different temporal scales and a second adversarial module is used to learn the transformation from one emotion to another. The first module is composed of a convolution layer with wavelet kernels so that the various temporal scales of F0 variations can be efficiently encoded. The single decomposition/transformation network allows to learn in an end-to-end manner the F0 decomposition that are optimal with respect to the transformation, directly from the raw F0 signal.

I. INTRODUCTION

Fundamental frequency (F0) is an essential acoustic feature in human speech communication and human-human interactions. As a key feature of speech prosody, F0 plays an important role in every aspect of speech communication: it conveys linguistic information (F0 helps to clarify the syntactic structure of an utterance or is used for semantic emphasis), para-linguistic information such as emotion or social attitude, and is even a part of the speaker identity through his speaking style. Accordingly, generative F0 modeling can be extremely useful in the fields of text-to-speech [1], voice identity conversion (VC) [2], and expressive voice conversion [3], by allowing a direct and parametric control of F0 to manipulate the expressivity of a voice (such as speaking style or emotions). By nature, F0 variations occur over different temporal scales each associated with specific functions, ranging from micro-variations to macro-contours such as accentuation, emotions, and modalities. To cover these specificities of F0 modelling, stylization methods [4], [5] and multi-level modelling [6]–[8] have been proposed.

Notably in VC, generative models such as Gaussian mixture models [9] and LSTM-based Sequence-to-Sequence models [10] were used to learn F0 transformations from neutral to expressive speech. Lately, various works focused on the use of Continuous Wavelet Transform (CWT), as an intermediary representation of the F0, on which Generative Adversarial Networks (GAN) models such as Dual-GAN [11], CycleGAN [12], VAW-GAN [13] or VA-GAN [14] are trained to learn transformations. A majority of those models is learnt on parallel data and by emotion pairs, which allows to learn a direct mapping between two different emotional versions of

an utterance while preserving a fixed and controlled linguistic content.

A promising approach called CWT Adaptive Scales (CWT-AS) was proposed by Luo et al. [11]. The CWT computes a decomposition of the F0 signal over wavelet kernels which allows a representation of F0 over different temporal scales [15], with various application in expressive voice conversion [3], [11], [13], [15], [16]. F0 modelling with CWT has been specified more recently upgraded with the possibility to compute the decomposition on arbitrary linguistic scales (e.g., phoneme, syllable, word, and utterance as described in [3]). An Adaptive-Scale (AS) algorithm [11] is described to select an optimal CWT representation for each pair of emotions, by selecting the scales that maximize in average the distance between the emotions in the CWT space.

From these selected scales, the CWT decomposition of the F0 contours is computed. Finally, the transformation function between each pair of emotion is learned from those representation using a Dual-GAN. Though this approach appears promising, it suffers from two main limitations: 1) The scale selection is only based on the maximization of the distance between the emotions, but ignores their reconstruction ability of the F0 signal. This may lead to poor F0 reconstruction which in turn would degrade the quality and the naturalness of the transformation; 2) The CWT-AS decomposition of the F0 signal and the dual-GAN are optimized independently which constitutes a bottleneck for training. Consequently, the CWT decomposition may not be optimal in the sense of the dual-GAN objective.

To overcome those limitations, we propose an end-to-end architecture to learn efficiently F0 transformation between emotions. The proposed neural architecture brings together the F0 decomposition and the dual-GAN into a single network, so that the CWT decomposition is optimized in the sense of the dual-GAN objective, and combining separation and reconstruction losses of the resulting decomposition. An application to the voice conversion of social attitudes shows that the proposed approach significantly improves the quality of the transformation by comparison with the CWT-AS approach.

II. PROPOSED METHOD

In this section, we introduce our proposal based on CWT-AS [11] and show how it differs by integrating the F0 stylization part on top of the transformation learning process which we refer to as end-to-end method for voice F0 conversion in II-A. Our contributions' concepts and technical details are given in II-B and II-C.

A. Framework overview

As our proposed VC system requires parallel data, sets of utterances X^a and X^b respectively relative to expressivity a and b are considered. A pair of utterances is then sampled and F0 sequences, source \mathbf{x}^a and target \mathbf{x}^b , are extracted. Aside from the expressivity, each utterance in a pair has the same content (linguistic content, speaker identity). The source and target F0 are given to what we called a Wavelet Kernel Convolutional Encoder (WKCE) denoted W_e . A classifier, denoted C , whose objective is to predict the expressivity is fed with WKCE outputs. As shown in Figure 1, these two modules must be seen as a pre-network (pN) for Dual-GAN (DG) that can be pre-trained as well as trained along with Dual-GAN forming an end-to-end system for f0 conversion.

B. Wavelet Kernel Convolutional Autoencoder

As a multiscale modelling method, CWT is entirely fitting when trying to represent both long and short-term dependencies, prosody is influenced by. As CWT can only be applied to continuous functions, a simple linear interpolation between voiced F0 segments is needed to obtain a continuous phrase-related F0 function which can then be sampled in a vector $\mathbf{x} \in [0, 1]^T$.

Our WKCE performs convolutions between the F0 signal \mathbf{x} and a wavelet kernel based on a mother wavelet $\psi_s \in \mathbb{R}^T$ defined for a time vector $\mathbf{t} \in \mathbb{R}^T$ as

$$\psi_s = \frac{2\pi^{-\frac{1}{4}}}{\sqrt{3}} \left(1 - \left(\frac{\mathbf{t}}{s}\right)^2\right) e^{-\frac{1}{2}\left(\frac{\mathbf{t}}{s}\right)^2} \quad (1)$$

Considering a wavelet kernel depending on N learnable parameters s that control the width of each wavelet composing the kernel, the contribution \mathbf{h}_x^s of the temporal level s to the F0 signal \mathbf{x} is the convolution between \mathbf{x} and ψ_s . Therefore, feed with \mathbf{x} , our WKCE module will output $W_e(\mathbf{x}) = [\mathbf{h}_x^{s_0}, \dots, \mathbf{h}_x^{s_N}] \in \mathbb{R}^{N \times T}$. If we denote W_r the reconstruction operation, then the reconstructed signal $\hat{\mathbf{x}}$ is given by

$$\hat{\mathbf{x}} = W_r(W_e(\mathbf{x})) = \frac{d_j \sqrt{d_t}}{C_d Y_0} \sum_{i=0}^{N-1} \mathbf{h}_x^{s_i} + \bar{\mathbf{x}} \quad (2)$$

with $\bar{\mathbf{x}}$ the average of \mathbf{x} , $d_j = 1.2$, $d_t = 0.125$, $C_d = 3.541$ and $Y_0 = 0.867$ (for details, see [17]).

If we denote \mathbb{E} , the mathematical expectation and consider \mathbf{x}^a and \mathbf{x}^b sampled from source and target distributions $P(\mathbf{x}^a)$ and $P(\mathbf{x}^b)$ respectively, this module can be trained for

reconstruction objective with respect to L_1 loss formulated as follows

$$L_{rec} = \mathbb{E}_{\mathbf{x}^a \sim P(\mathbf{x}^a)} (\|W_r(W_e(\mathbf{x}^a)) - \mathbf{x}^a\|_1) + \mathbb{E}_{\mathbf{x}^b \sim P(\mathbf{x}^b)} (\|W_r(W_e(\mathbf{x}^b)) - \mathbf{x}^b\|_1) \quad (3)$$

A constraint of classification on the CWTs latent space can be added, W_e and C are trained with respect to L_{cl} , the cross-entropy (CE) loss between the predicted source expressivity $\hat{a} = C(W_e(\mathbf{x}^a))$ and the true value a summed with the CE between \hat{b} and b .

$$L_{cl} = \mathbb{E}_{\mathbf{x}^a \sim P(\mathbf{x}^a)} [a * C(W_e(\mathbf{x}^a))] + \mathbb{E}_{\mathbf{x}^b \sim P(\mathbf{x}^b)} [b * C(W_e(\mathbf{x}^b))] + \mathbb{E}_{\mathbf{x}^b \sim P(\mathbf{x}^b)} (1 - a) [1 - \log(C(W_e(\mathbf{x}^a)))] + \mathbb{E}_{\mathbf{x}^a \sim P(\mathbf{x}^a)} (1 - b) [1 - \log(C(W_e(\mathbf{x}^b)))] \quad (4)$$

C. Model

In this paper we focus on a specific GAN network called Dual-GAN which is capable of learning a mapping between parallel pairs of data. This network is based on two concepts: 1) Adversarial learning [18], which is to train a generative model to find a solution in a min-max game between two neural networks, called as generator G and discriminator D . 2) Dual supervised learning [19] which is to train the models of two dual tasks simultaneously exploiting the probabilistic correlation between them to regularize the training process. Combining those breakthroughs allows to take advantage of the GAN's ability to produce realistic transformations as well as the significant improvements due to dual supervised learning.

This second point implies that both forward and inverse transformations, respectively $G_{a \rightarrow b} : (W_e(\mathbf{x}^a), z^a) \rightarrow \mathbf{x}^b$ and $G_{b \rightarrow a} : (W_e(\mathbf{x}^b), z^b) \rightarrow \mathbf{x}^a$, are learned jointly, where z^a and z^b are random independant noises provided in the form of dropout at each layer of G_a and G_b . A first loss $L_{a \leftrightarrow b}$ is required to train $G_{a \rightarrow b}$, $G_{b \rightarrow a}$ and W_e .

$$L_{a \leftrightarrow b} = \mathbb{E}_{(\mathbf{x}^a, \mathbf{x}^b) \sim P(\mathbf{x}^a, \mathbf{x}^b)} (\|W_r(G_{a \rightarrow b}(W_e(\mathbf{x}^a))) - \mathbf{x}^b\|_1) + \mathbb{E}_{(\mathbf{x}^a, \mathbf{x}^b) \sim P(\mathbf{x}^a, \mathbf{x}^b)} (\|W_r(G_{b \rightarrow a}(W_e(\mathbf{x}^b))) - \mathbf{x}^a\|_1) \quad (5)$$

In the same time, D_a discriminates between converted outputs \hat{X}^b of $G_{a \rightarrow b}$ and real samples of domain X^b , D_b does analogously the same to complete the adversarial mechanism. The adversarial loss L_{ADV} is required to train $G_{a \rightarrow b}$, $G_{b \rightarrow a}$, D_a , D_b and W_e

$$L_{adv} = \mathbb{E}_{\mathbf{x}^a \sim P(\mathbf{x}^a)} [D_a(W_e(\mathbf{x}^a))] + \mathbb{E}_{\mathbf{x}^b \sim P(\mathbf{x}^b)} [D_b(W_e(\mathbf{x}^b))] + \mathbb{E}_{\mathbf{x}^b \sim P(\mathbf{x}^b)} [1 - \log(D_a(G_{b \rightarrow a}(W_e(\mathbf{x}^b)))] + \mathbb{E}_{\mathbf{x}^a \sim P(\mathbf{x}^a)} [1 - \log(D_b(G_{a \rightarrow b}(W_e(\mathbf{x}^a)))] \quad (6)$$

A third constraint called Dual loss is added so as to strengthen the intrinsic connection between $G_{a \rightarrow b}$ and $G_{b \rightarrow a}$, it can be understood as a regularization of the process.

$$L_{dual} = \mathbb{E}_{(\mathbf{x}^a, \mathbf{x}^b) \sim P(\mathbf{x}^a, \mathbf{x}^b)} (\|W_e(\mathbf{x}^a) * G_{a \rightarrow b}(W_e(\mathbf{x}^a)) - W_e(\mathbf{x}^b) * G_{b \rightarrow a}(W_e(\mathbf{x}^b))\|_1) \quad (7)$$




Fig. 1. End-to-end neural architecture for F0 voice conversion. On the right, F0 decomposition over 4 of the learned scales obtained for the source (red) and target (green) expressivities.

Therefore two final losses can be formulated for pre-Net pretraining and proper Dual-GAN training, respectively L_{pN} and L_{DG} with α , β , λ and γ respectively weighting reconstruction, classification, transformation and dual objectives.

$$L_{pN} = \alpha L_{rec} + \beta L_{cl} \quad (8)$$

$$L_{DG} = \lambda L_{a \leftrightarrow b} + L_{adv} + \gamma L_{dual} \quad (9)$$

III. EXPERIMENTS

A. Dataset

For our experiments we used the freely available speech database Att-HACK [20]. The database comprises 25 speakers interpreting 100 utterances in 4 social attitudes : friendly, distant, dominant and seductive, later denoted FR, DIST, DOM and SED respectively. With 3 to 5 repetitions each per attitude for a total of around 30 hours of speech, the database offers a wide variety of prosodic strategies in the expression of attitudes. Prosodic features conveying expressivity has been shown to be speaker dependent [21], for this reason, two speakers were selected and used independently for both training and validation : a female (F08) and a male (M07), representing almost 400 utterances each. The *train/valid* split was 80/20 % and has been done linguistically.

B. Implementation details

1) *Input pipeline*: We extracted fundamental frequency from the speech signal by using SWIPEP algorithm [22]. All F0 sequences are sampled to 1ms (as recommended in [11]), passed to $\log(F0)$ and a linear interpolation has been processed between voiced segments. For each pair, a mapping between syllables starting and ending times has been done to align source and target. Once pairs are aligned syllable-wise, F0 sequences are padded with zeros up to a value $T = 4000$ corresponding to the longest sentence in the dataset.

2) *Architecture design*: Our WKCE denoted W_e , with $N = 32$ learnable scales, has been implemented as a custom layer. A constraint of growth has been added on the range of scales to ensure the continuity of the learned CWTs. The output of

W_e of shape $[32, 4000]$ is unpadded and temporally sliced to form batches of shape $[\text{batch_size}, 32, 512]$.

The classifier is built using convolutional blocks composed of three convolutional layers starting from 32 up to 128 filters. Each block uses 3×3 convolutions with *ReLU* activation, a dropout of 0.2, padding mode same and pooling operations using both strides and max pooling 2D with values 2 and 4, respectively, reducing features and time. Those blocks are followed with a flatten layer and two dense layers with respectively, 1000 units and a *ReLU* activation, and 2 units and a *softmax* activation.

Two configurations of our pre-network pN can be distinguished :

- *config_A* : $pN = \{W_e\}$ learns the CWT scales regarding the CWT reconstruction objective ($\alpha = 1, \beta = 0$)
- *config_B* : $pN = \{W_e + C\}$ learns the CWT scales regarding both the CWT reconstruction and the CWT related attitude classification objectives ($\alpha = 10, \beta = 1$)

The architecture of the Dual-GAN itself as well as the contribution of each module, $G_{a \rightarrow b}$, $G_{b \rightarrow a}$, D_a and D_b , in training process are taken from [11].

3) *Training procedure*: W_e is being fed with batches of size 1 so as to allow the Dual-GAN to process on phrase related batches. pN is pre-trained minimizing L_{pN} depending on the considered configuration. CWTs are learned with respect to voiced segments by using a voicing binary mask of 1 (voiced) and 0 (unvoiced). Then $\{DG + W_e\}$ is trained minimizing L_{DG} with $\lambda = 5$ and $\gamma = 15$. For *config_B*, the classification scores after pre-training were passed as sample weights for Dual-GAN training. ADAM optimizer with 0.0001 as learning rate has been used. All codes are written in Python-Tensorflow 2.1, the baseline has been re-implemented.

IV. RESULTS AND DISCUSSIONS

A. Objective Evaluation

Considering 4 attitudes, we evaluated 12 transformations (6 forward, 6 invert), the overall results are shown in Table I,

using the RMSE between 1) the original F0 and the reconstruction obtained from its representation, 2) the converted F0 and the corresponding target F0. An example of conversion is depicted in Figure 2.

First of all, we observe that the proposed end-to-end system for F0 voice conversion outperforms the traditional *baseline*, in average over all categories, for both configurations. Secondly our most sophisticated configuration *config_B* achieves slightly better results than *config_A*.

Models	RMSE (Hz)	
	Reconstruction	Transformation
<i>baseline</i>	17.32	21.71
<i>config_A</i>	9.16	19.15
<i>config_B</i>	13.4	18.83

TABLE I

A COMPARISON OF THE RMSE RESULTS OF THE *baseline*, *config_A* AND *config_B* FOR RECONSTRUCTION AND TRANSFORMATION




Fig. 2. Example of F0 transformation from distant to dominant for speaker M07 for *baseline* and ours *config_A* and *config_B*.

B. Subjective Evaluation

We further conducted a listening experiment to compare the two proposed configurations with the baseline CWT-AS in terms of attitude similarity. We perform XAB test to assess the emotion similarity by asking listeners to choose between two converted utterances (*baseline* and one of our configurations) the one which sounds more similar to the original target. The overall results are reported in Figure 3. Our proposed system in *config_B* outperforms the baseline for each of the considered transformations. The results for *config_A*, slightly lower, still outperforms the *baseline*.




Fig. 3. The XAB preference results with 95% confidence interval between the *baseline* and ours *config_A* and *config_B* regarding attitude similarity.

C. Scales distributions

This part provides an a posteriori comparison of the F0 scales distribution as obtained by the description presented

in parts IV-A and IV-B. Each transformation (forward and backward) between a pair of attitudes is associated with a set of temporal scales that are used to compute the CWT representations used for the conversion. Consequently, each transformation can be described by a distribution of the temporal scales that are used to convert the F0 optimally. Figure 4 presents the distribution selected by the baseline CWT-AS algorithm and learned by our proposed contribution, as obtained for speaker F08 for the six pairs of attitudes.

The best performance being obtained with the *config_B* of the proposed contribution, we further investigated by comparison of the F0 scales distribution. First, one can clearly observe that the temporal distribution of the *config_B* is wider than the others, the transformation covering a wide range of temporal scales from the micro variations over the phonemes to the global contours of the sentence. Additionally, the distributions associated with the *baseline* and the *config_A* appear mostly independent with respect to the transformation pair, while the distribution associated with the *config_B* tend to be more varied depending on the transformation pair. This suggests that the *config_B* may adapt more efficiently to the specificities of each pair




Fig. 4. CWT scale distributions for the three considered algorithms: *baseline*, *config_A* and *config_B*. For interpretation, P, SY, W and SE markers respectively denote average duration of the phone, syllable, word and sentence linguistic units.

V. CONCLUSION

In this paper, we propose a end-to-end framework for F0 transformation in the context of expressive voice conversion, bringing together the F0 decomposition in different temporal levels and its transformation in a single network. Both objective and subjective evaluations showed our method can achieve better performance than the baseline. We aim at generalizing for multi-speaker F0 conversion and to avoid pair-learning by building an expressive embedding. An online page featuring conversion examples is available at http://recherche.ircam.fr/anasy/VC_demo/index.html.

ACKNOWLEDGEMENTS

This research is supported by the MoVe project: “MOdelling of speech attitudes and application to an expressive conversational agent”, and funded by the Paris Region Ph2D grant.

REFERENCES

- [1] Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ Skerry-Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Fei Ren, Ye Jia, and Rif A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," 2018.
- [2] T. Kaneko and H. Kameoka, "CycleGAN-vc: Non-parallel voice conversion using cycle-consistent adversarial networks," in *2018 26th European Signal Processing Conference (EUSIPCO)*, 2018, pp. 2100–2104.
- [3] Zhaojie Luo, Jinhui Chen, Tetsuya Takiguchi, and Yasuo Arikawa, "Emotional voice conversion using neural networks with arbitrary scales f0 based on wavelet transform," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2017, no. 1, pp. 18, 2017.
- [4] Jonathan Teutenberg, Catherine Watson, and Patricia Riddle, "Modelling and Synthesising F0 contours with the Discrete Cosine Transform," in *International Conference on Acoustics, Speech, and Signal Processing*, Las Vegas, U.S.A., 2008, pp. 3973–3976.
- [5] Nicolas Obin and Julie Béli o, "Sparse coding of pitch contours with deep auto-encoders," in *International Conference on Speech Prosody*, 2018, pp. 799–803.
- [6] Nicolas Obin, Anne Lacheret, and Xavier Rodet, "Stylization and Trajectory Modelling of Short and Long Term Speech Prosody Variations," in *Interspeech*, Florence, Italy, 2011, pp. 2029–2032.
- [7] Xiang Yin, Ming Lei, Yao Qian, Frank K. Soong, Lei He, Zhen-Hua Ling, and Li-Rong Dai, "Modeling F0 trajectories in hierarchically structured deep neural networks," *Speech Communication*, vol. 76, pp. 82–92, 2016.
- [8] Xin Wang, Shinji Takaki, and Junichi Yamagishi, "An RNN-Based quantized F0 model with Multi-Tier feedback links for Text-to-Speech synthesis," in *Interspeech*, 2017.
- [9] C. Veaux and X. Rodet, "Intonation conversion from neutral to expressive speech," in *Interspeech*, 2011.
- [10] Carl Robinson, Nicolas Obin, and Axel Roebel, "Sequence-to-Sequence Modelling OF F0 For Speech Emotion Conversion," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Brighton, United Kingdom, May 2019.
- [11] Z. Luo, J. Chen, T. Takiguchi, and Y. Arikawa, "Emotional voice conversion using dual supervised adversarial networks with continuous wavelet transform f0 features," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 10, pp. 1535–1548, 2019.
- [12] Kun Zhou, Berrak Sisman, and Haizhou Li, "Transforming spectrum and prosody for emotional voice conversion with non-parallel training data," 2020.
- [13] Kun Zhou, Berrak Sisman, Mingyang Zhang, and Haizhou Li, "Converting anyone's emotion: Towards speaker-independent emotional voice conversion," 2020.
- [14] Zhaojie Luo, Jinhui Chen, Tetsuya Takiguchi, and Yasuo Arikawa, "Neutral-to-emotional voice conversion with cross-wavelet transform f0 using generative adversarial networks," *APSIPA Transactions on Signal and Information Processing*, vol. 8, pp. e10, 2019.
- [15] H. Ming, D. Huang, M. Dong, H. Li, L. Xie, and S. Zhang, "Fundamental frequency modeling using wavelets for emotional voice conversion," in *Affective Computing Intell. Interact.*, 2015, p. 804–809.
- [16] Huaiping Ming, Dongyan Huang, Lei Xie, Jie Wu, Minghui Dong, and Haizhou Li, "Deep bidirectional LSTM modeling of timbre and prosody for emotional voice conversion," in *Interspeech 2016*, 2016.
- [17] C. Torrence and G. P. Compo, "A practical guide to wavelet analysis," *Bulletin of the American Meteorological Society*, vol. 79, pp. 61–78, 1998.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds., pp. 2672–2680. Curran Associates, Inc., 2014.
- [19] Y. Xia, T. Qin, W. Chen, J. Bian, N. Yu, and T.-Y. Liu, "Dual supervised learning," 2017.
- [20] Cl ment Le Moine and Nicolas Obin, "Att-HACK: An Expressive Speech Database with Social Attitudes," in *Speech Prosody*, Tokyo, Japan, May 2020.
- [21] Berrak Sisman and Haizhou Li, "Wavelet analysis of speaker dependent and independent prosody for voice conversion," 09 2018.
- [22] Arturo Camacho, *SWIPE: A Sawtooth Waveform Inspired Pitch Estimator for Speech and Music*, Ph.D. thesis, University of Florida, 2007.