

# TG-CRITIC: A TIMBRE-GUIDED MODEL FOR REFERENCE-INDEPENDENT SINGING EVALUATION

Xiaoheng Sun<sup>1,\*</sup>, Yuejie Gao<sup>1,\*</sup>, Hanyao Lin<sup>2,\*</sup> and Huaping Liu<sup>1,†</sup>

<sup>1</sup> NetEase Cloud Music, Shanghai, China

<sup>2</sup> Fudan University, Shanghai, China

## ABSTRACT

Automatic singing evaluation independent of reference melody is a challenging task due to its subjective and multi-dimensional nature. As an essential attribute of singing voices, vocal timbre has a non-negligible effect and influence on human perception of singing quality. However, no research has been done to include timbre information explicitly in singing evaluation models. In this paper, a data-driven model TG-Critic is proposed to introduce timbre embeddings as one of the model inputs to guide the evaluation of singing quality. The trunk structure of TG-Critic is designed as a multi-scale network to summarize the contextual information from constant-Q transform features in a high-resolution way. Furthermore, an automatic annotation method is designed to construct a large three-class singing evaluation dataset with low human-effort. The experimental results show that the proposed model outperforms the existing state-of-the-art models in most cases.

**Index Terms**— Singing evaluation, timbre embedding, high-resolution network

## 1. INTRODUCTION

Automatic singing evaluation aims to assess the quality of singing performances without the participation of music experts, thus reducing labor costs. The applications of the task include the distribution of singing content on the Internet and the discovery of musicians. With the popularity of online music services and karaoke singing applications, there is a growing demand for more advanced evaluation systems for singing voices, making automatic singing evaluation a hot research topic in recent years.

Because singing evaluation by humans is empirically based and listener dependent, it is difficult to use automatic systems in substitution of human experts. However, previous studies have shown that humans rely more on some common, objective features than subjective preferences when judging the quality of singing [1]. These features such as intonation accuracy, tonal stability, and rhythm consistency can be detected automatically, which makes computational techniques feasible for this task. Also, the criteria of singing evaluation are inevitably multi-dimensional. A reliable singing evaluation system can be established only by observing objective metrics of a singing work in a similar way as music experts do.

Depending on whether a reference melody is required, the existing automatic singing evaluation systems can be roughly divided into two types: reference-dependent approaches and reference-independent ones. A typical reference-dependent approach needs one ideal reference, such as an excellent singing recording or the correct melodic line. The more similar the singing is to the reference,

the higher the score or the ranking level [2]. Most studies focus on pitch accuracy [3, 4, 5], while some studies design hand-crafted audio features to obtain scores from other dimensions such as dynamic [2, 6], rhythm [2, 6], voice quality [6], timbre brightness [7, 8], enthusiasm [9], vibrato [6], tremolo [10], *etc.* There are also works using metric learning to map singing voices to embeddings for similarity comparison [11, 12]. The advantages of reference-dependent approaches are high interpretability and low complexity, while the disadvantages are as below: **a)** Manual preparation/proofreading of references is needed; **b)** It encourages the imitation of references rather than personalized interpretation.

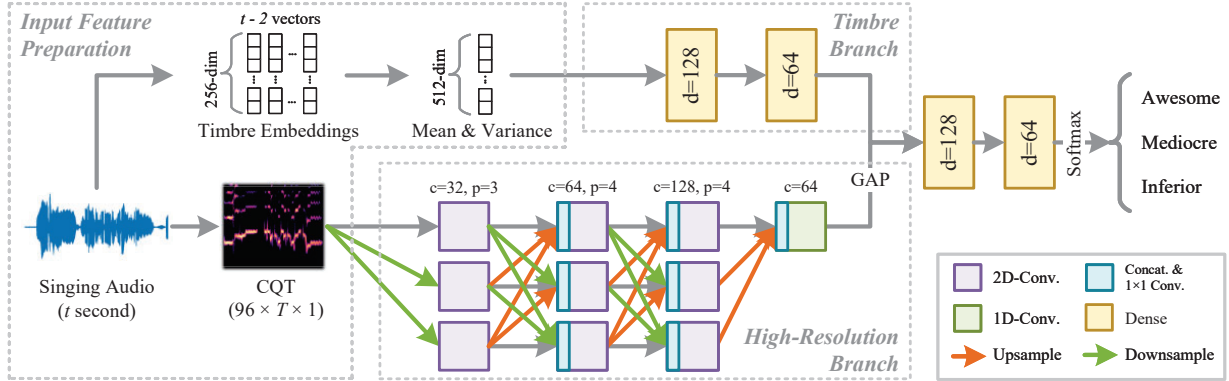
For human experts, studies have shown that even if the melody is new to them, they can make a highly consistent evaluation [1], indicating the feasibility of the reference-independent approaches. The traditional reference-independent singing evaluation mainly focused on intonation characteristics [13, 14]. Since Zhang *et al.* [15] proposed a convolutional neural network (CNN) based model trained on two-class data, the data-driven methods have been introduced into this field. Gupta *et al.* [16] adopted a convolutional recurrent neural network (CRNN) framework, where the inputs are Mel-spectrograms along with pitch histograms. Huang *et al.* [17] further compared Mel-spectrogram with different mid-level input features such as constant-Q transform (CQT) and Chromagram. Using a similar CRNN structure, Li *et al.* [18] trained multi-task models for both pitch and overall scores. Gupta *et al.* [19] transplanted the idea of pitch histogram to rhythm histogram for rhythm evaluation. The reference-independent methods are truly end-to-end, which helps to further reduce the need for human-effort and expand application scenarios.

Although great performance has been achieved by previous studies, we notice that so far there is no such method that involves timbre information explicitly in the model. As previous studies [4, 8, 20] have shown, the vocal timbre is an important attribute that affects human perception of singing. But unlike other attributes such as intonation and rhythm, timbre is a multifarious set of abstract sensory attributes [21] and is difficult to be extracted by models. Therefore, we believe it is helpful to involve timbre information explicitly rather than to have the model learn the features from scratch.

In this paper, we explore adding timbre embeddings as the model inputs and propose a *timbre-guided* singing evaluation model named TG-Critic. The trunk structure of TG-Critic is designed as a multi-scale CNN-based network to summarize the contextual information from the input CQT in a high-resolution way. In addition to the CQT, TG-Critic also takes a vector involving timbre information derived from the audio using a timbre embedding model. Furthermore, to alleviate the problem of insufficient data, we propose an iterative automatic annotation method and construct a large singing dataset YJ-16K. Experimental results show that by including timbre information explicitly, the model achieves higher accuracy and outperforms existing state-of-the-art methods.

\*These authors contributed equally to this research.

†Corresponding author. E-mail: liuhuaping@corp.netease.com



**Fig. 1.** The overall architecture of TG-Critic. ‘c’ denotes the channel number of the convolutional layer, while ‘p’ denotes the pooling length along the frequency axis. All the 2D convolutional layers use the kernel size of  $3 \times 3$ , while the 1D convolutional layer’s kernel length is 5. All the layers except the last dense layer are followed by the ELU activation function.

## 2. APPROACH

The overall architecture of TG-Critic is shown in Fig. 1. It has two branches: 1) a *High-Resolution Branch* mainly composed of convolutional layers processing the input CQT into high-level semantic representations for singing evaluation; 2) and a *Timbre Branch* processing the timbre vector derived by a timbre embedding model. In this section, we first introduce the timbre embedding method and Timbre Branch. Then, the structure of High-Resolution Branch is detailed. Finally, we describe the construction of the singing dataset YJ-16K.


### 2.1. Timbre Embedding & Timbre Branch

Timbre allows us to distinguish the difference between the two voices besides pitch, loudness, and duration [22]. To generate timbre representations that we can input into a deep learning model, we adopt the embedding model *CROSS* proposed by Lee and Nam [23] and use the pretrained model<sup>1</sup> in experiments.

*CROSS* is a metric learning-based embedding model designed for singer-relevant tasks, especially singer identification. A previous study has proved that singer identity mainly consists of two parts: timbre and singing style [24]. Therefore, the output vector should include enough information to distinguish vocal timbre and we believe it is suitable for timbre embedding.

The input of *CROSS* is a three-second long audio clip while the output is an embedding vector of 256 dimensions. In experiments, we first slice each audio into segments of three seconds with the step size as one second and then derive the embedding vectors for each segment. Therefore, for audio with a length of  $t$  seconds, a “timbregram” with a size of  $(t - 2, 256)$  will be obtained, which depicts the variation of vocal timbre in the audio. To alleviate the influence of loudness, we perform min-max normalization on the timbregram. Finally, the mean and variance along the time axis are calculated and then concatenated as a vector of length 512, which we used as the input of TG-Critic. We excluded invalid timbre embeddings (singing voice frames are less than 70%) in the calculation.

To prove that information related to the singing quality is involved in the derived vectors, we randomly choose some singing samples with three-level quality annotations “Awesome”, “Mediocre” and “Inferior”. Through the method mentioned above, from each



**Fig. 2.** The t-SNE visualization of the timbre vectors derived from singing samples with different qualities. M.S. and F.S. denote Male Singers and Female Singers respectively.

audio, a 512-dimension timbre vector can be obtained. Then these timbre vectors are visualized using t-Distributed Stochastic Neighbor Embedding (t-SNE) [25], as shown in Fig. 2. It shows that for male and female singing respectively, vectors of the same quality level are closer to each other than those of different levels. Furthermore, in an informal experiment, we use only these vectors to train a multilayer perceptron (MLP) classification model and get an accuracy of 62%. These results indicate that the timbre vectors contain information helpful for automatic singing evaluation.

As shown in Fig. 1, Timbre Branch with two dense layers is designed to further process the 512-dimension timbre vectors into representations suitable for the fusion with the outputs of High-Resolution Branch. Finally, we obtain an output vector of 64 dimensions.

### 2.2. Structure of High-Resolution Branch

We use CQT as the input mid-level feature for High-Resolution Branch because previous studies [17, 18, 19] have proved its effectiveness. To extract features, the original audio data is first resampled to 16 kHz. For CQT computation, we set the hop size to 512. For each frame, there are 96 bins and each octave is covered by 24 bins. The minimum frequency is set to  $E/2$  to better match the pitch range of singing voices. For training and evaluation, we set the input length of CQT as 256 frames (roughly 8 seconds). Note that the input audio of TG-Critic can be of arbitrary length.

<sup>1</sup><https://github.com/kyungyunlee/mono2mixed-singer>

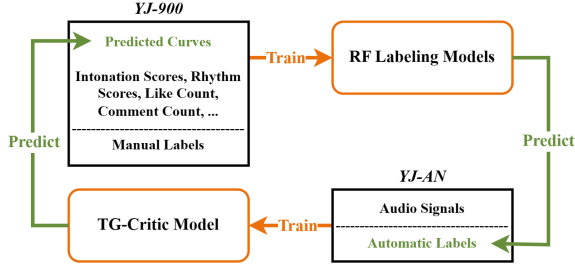


Fig. 3. The iterative automatic annotation of YJ-16K dataset.

For singing evaluation, long-term dependencies and local details are both indispensable. As many studies [15, 17, 18] did, we use a CNN-based structure as the backbone of our model to better detect local patterns. However, a few convolution layers cannot model long-distance dependence. On the other hand, the over-stacking of convolutional layers may lead to excessive parameters. Downsampling operations are usually introduced to expand the context range that convolution kernels can capture. But downsampling also causes a loss of detailed information. Besides, some studies [17, 18] adopted the idea of CRNN, applying a recurrent layer after CNN to capture long-term dependencies. However, the involvement of RNN undoubtedly reduces the model’s efficiency.

Inspired by the success of HRNet [26] designed for computer vision tasks, we introduce a similar multi-scale structure into the proposed model. In such a structure, downsampling is performed on the input features to obtain low-resolution representations to facilitate the extraction of global information. After the process of convolutional layers, the low-resolution features are upsampled and merged with high-resolution features. The procedure of downsampling and upsampling is repeated several times, and it helps aggregate contextual information from different scales. Because of the existence of high-resolution features, the detailed patterns are not degraded during the process. Specifically, the input CQT of  $T$  frames is first downsampled to  $T/2$  and  $T/4$  over the time axis (we assume that  $T$  is an integral multiple of 4). These features are then processed by  $3 \times 3$  convolutional layers. After average pooling over the frequency axis to summarize frequency bands, the features are rescaled and fused. Average pooling and nearest neighbor interpolation are used to perform rescaling. For the fusion operation, the concatenation is followed by a  $1 \times 1$  convolutional layer to exchange information and adjust the channel number. Meanwhile, on the frequency axis the features are gradually downsampled. After three multi-scale processes, the features are fused again. We reshape the obtained feature to  $(T, 768)$  and use a 1D convolutional layer to adjust the channel numbers to 64. After global average pooling, a vector of 64 dimensions is obtained.

After obtaining the output vectors from the two branches, we concatenate them and send the fused vector into a dense layer with exponential linear unit (ELU) activation [27] to exchange information. Finally, we apply a dense layer with softmax activation to produce the classification result for singing quality.

### 2.3. Construction of YJ-16K Dataset

We construct a three-class singing dataset named YJ-16K. The dataset contains 32,623 users’ unaccompanied singing pieces from YinJie<sup>2</sup>. We collect metadata from the application including scoring

<sup>2</sup><https://k.163.com/> is a karaoke application.

		YJ-16K	Total	A	M	I	M.S.	F.S.
YJ-900	test	894	296	302	296	430	464	
YJ-AN	total	15000	5000	5000	5000	7683	7317	
	train	14100	4700	4700	4700	7203	6897	
	valid	900	300	300	300	480	420	

Table 1. The number of samples in YJ-16K dataset. M.S. denotes Male Singers while F.S. denotes Female Singers.

data such as intonation scores, rhythm scores, *etc.*, and user behavior data such as the numbers of likes and comments for each song. The dataset contains a variety of singing from poor to professional. We aim to label these samples into three classes: Awesome (A) - works with excellent execution and impressive voice quality; Mediocre (M) - works with good intonation, but with ordinary expressiveness; Inferior (I) - works with weak intonation or poor expressiveness.

To reduce labor costs of labeling, we train models for automatic annotation. First, a subset of 894 samples named YJ-900 is separated and manually annotated by music experts. Using YJ-900’s metadata and manual labels, we train two random forest (RF) binary classification models to separate Awesome and Inferior samples respectively. Then, the RF models are used to annotate the rest of the original dataset, called YJ-AN. However, these automatic labels are not reliable enough, because the metadata we obtained was affected by various external factors. For example, covers of hit songs often tend to get higher user ratings.

To improve the reliability of annotations, we design an “iterative automatic annotation” method as shown in Fig. 3. We first train the TG-Critic model using the audio signals and automatic labels of YJ-AN. By this model, a predicted “evaluation curve” is obtained for each audio sample of YJ-900 (as TG-Critic’s inputs are 8-second segments), which is then used as a part of the inputs of the RF Models in the next iteration. The iteration repeats four times, and the detail of the finally obtained dataset is shown in Table 1. The effects of iterative automatic annotation can be seen in Section 3.3.

## 3. EVALUATION

### 3.1. Ablation Studies

To demonstrate how much the timbre vectors contribute to the model, we compare the performance of the models with and without Timbre Branch. For the structure with Timbre Branch (the proposed TG-Critic), two models are trained by different training strategies: **a**) *TG-Critic-1S*: The High-Resolution Branch and the Timbre-Branch are trained together in one step; **b**) *TG-Critic-2S*: A two-step training strategy is applied where the High-Resolution Branch is first trained and frozen, and then the Timbre Branch is trained using smaller learning rate and fewer epochs. We also trained a model *CQT-Only* with only High-Resolution Branch and CQT input for comparison. Among the models mentioned above, *TG-Critic-1S* and *CQT-Only* are both trained by 200 epochs with a learning rate of 0.0001. *TG-Critic-2S* is trained 100 epochs for each step, and the learning rate is set as 0.0001 in the first step and 0.00005 in the second step. For all three models, the Adam optimizer is applied in training to minimize categorical cross-entropy, and one epoch consists of 500 batches where the batch size is set as 32.

As shown in Table 2, the model without Timbre Branch performs worse than the models with Timbre Branch. Among the three classes, *CQT-Only* shows the highest recall and the lowest precision

Model	Precision (%)			Recall (%)			Acc. (%)
	A	M	I	A	M	I	
<i>TG-Critic-1S</i>	83.5	71.6	84.0	90.5	69.2	79.7	79.8
<i>TG-Critic-2S</i>	<b>87.2</b>	<b>73.6</b>	86.7	89.9	<b>75.5</b>	81.8	<b>82.3</b>
<i>CQT-Only</i>	84.3	69.8	79.5	88.9	63.6	<b>82.4</b>	78.2
<i>TG-Simple</i>	82.1	68.4	<b>88.7</b>	<b>92.9</b>	72.5	71.6	79.0

**Table 2.** Results of ablation studies on YJ-900 dataset.

Model	Param.	YJ-900	PESnQ-DS		NUS48E	
		Acc.	Acc.	Corr.	Acc.	Corr.
<i>Kuaishou</i> [15]	1.97M	68.3	85.0	0.858	68.8	0.497
<i>NUS20</i> [17]	0.72M	76.3	85.0	0.930	68.8	0.552
<i>NUS21</i> [18]	1.45M	78.4	85.0	0.925	72.9	0.548
<i>TG-Critic-1S</i>	0.82M	79.8	80.0	0.927	72.9	<b>0.671</b>
<i>TG-Critic-2S</i>		<b>82.3</b>	<b>95.0</b>	<b>0.933</b>	<b>77.1</b>	0.631

**Table 3.** Results of the proposed and baseline models on YJ-900, PESnQ-DS, NUS48E datasets. The accuracy values are percentiles. (Acc. = Accuracy, Corr. = Pearson Correlation Coefficient.)

on the inferior singings, indicating the model without explicit timbre information is inclined to underestimate the singing quality. For the two models with Timbre Branch, *TG-Critic-2S* using the two-step training strategy presents a significant enhancement in performance. We believe this is mainly because Timbre Branch is a much simpler structure than High-Resolution Branch, and it is difficult for the two branches to converge simultaneously in one-step training.


Furthermore, to evaluate the idea of using a “high-resolution” network, we remove the High-Resolution Branch in the model named *TG-Simple*. Instead, a simple CNN structure stacked with three  $3 \times 3$  convolutional layers is used. We carefully adjust the filter numbers in the layers to remain a similar number of parameters, and the same training strategy as *TG-Critic-2S* is used to ensure the model’s best performance. As shown in Table 2, the accuracy decreases by 3.3% compared to *TG-Critic-2S*, proving that the high-resolution structure helps to produce better representations for singing evaluation.

### 3.2. Comparison with Previous Works

To compare with previous works, we reproduce three baseline models: *Kuaishou* [15] (Bi-DenseNet-27 model), *NUS20* [17] (CPH-CRNN model), and *NUS21* [18] (Framework 3). All three models are reference-independent singing evaluation models based on deep learning. We follow the hyperparameters described in the papers to ensure the models’ performances on our training dataset YJ-AN.

In addition to YJ-900, we use two public datasets PESnQ-DS [6] and NUS48E [28] for tests. Recordings of PESnQ-DS (two songs, each performed by 10 singers) cover singing ability from poor to professional and they are annotated with overall scores, pitch scores, and rhythm scores. NUS48E contains 48 samples (20 songs) performed by 12 singers between mediocre and good level, without quality annotations. Besides, three of the PESnQ-DS recordings are from the NUS48E dataset.

To collect subjective ratings on NUS48E, we recruit 12 human judges to give an overall singing quality score out of 5 to each recording compared to the reference singing samples. We also conduct the same rating procedure on PESnQ-DS and find our scores are highly consistent with the overall scores provided by the authors, achieving



**Fig. 4.** Performance of models trained on datasets of different annotation iterations.

the Pearson correlation coefficient (PCC) of 0.91, which proves the reliability of our annotations.

Considering the singers’ singing ability of NUS48E is from mediocre to good level, we divide the NUS48E samples into classes **A** and **M** using a threshold (resulting in 29 as **A** and 19 as **M**). On PESnQ-DS, we calculate the average scores of all annotators (12 by us and 5 by [6]) for comparison, and divide samples into classes **A**, **M**, and **I** by two thresholds (resulting in 8 as **A**, 7 as **M**, and 5 as **I**).

The experimental results are shown in Table 3. Compared with other models, *TG-Critic-2S* achieves the highest accuracy for all three datasets. To make a comprehensive comparison, we obtain a weighted score for each prediction using the output probability distribution. We annotate predicted probabilities for class **A**, **M**, and **I** be  $P_A$ ,  $P_M$ , and  $P_I$  respectively. Then the weighted score is calculated as:  $P_A \times 1.0 + P_M \times 0.5 + P_I \times 0.0$ , ranging within  $[0, 1]$ . The PCCs between the weighted scores and the ground truth are shown in Table 3. The high correlations indicate that these scores from probability distribution are capable of measuring the singing quality. *TG-Critic-1S* and *TG-Critic-2S* show better performance than the baseline models, clearly confirming the effectiveness and robustness of our proposed model<sup>3</sup>.

### 3.3. Effects of Iterative Automatic Annotation

To demonstrate the effects of iterative automatic annotation, we train models with datasets after each annotation iteration using the two-step training strategy mentioned in Section 3.1. Then, the models are evaluated on the manually annotated YJ-900.

As shown in Fig. 4, the model’s performance improves significantly after each iteration, especially for the first three ones, which indicates that the annotations are increasingly closer to human assessment. The results prove that the iterative annotation process is effective in producing reliable annotations.

## 4. CONCLUSION

In this paper, we have proposed TG-Critic, a timbre-guided singing evaluation model independent of the reference melody. The proposed model includes timbre information explicitly by using timbre embedding as one of the model inputs. A multi-scale structure is introduced to process the CQT features in a high-resolution way. We also construct a large singing dataset YJ-16K with annotations labeled by an iterative automatic annotation method. Experimental results show the proposed model outperforms the existing state-of-the-art models in most cases.

<sup>3</sup>The detail annotations for PESnQ-DS and NUS48E and further experimental results are available at <https://github.com/YuejieGao/TG-CRITIC>

## 5. REFERENCES

- [1] Tomoyasu Nakano, Masataka Goto, and Yuzuru Hiraga, "Subjective evaluation of common singing skills using the rank ordering method," in *Ninth International Conference on Music Perception and Cognition*. Citeseer, 2006, pp. 1507–1512.
- [2] Wei-Ho Tsai and Hsin-Chieh Lee, "An automated singing evaluation method for karaoke systems," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2011, pp. 2428–2431.
- [3] Partha Lal, "A comparison of singing evaluation algorithms," in *Proceedings of INTERSPEECH 2006*, 2006, vol. 5, pp. 2298–2301.
- [4] Chuan Cao, Ming Li, Jian Liu, and Yonghong Yan, "A study on singing performance evaluation criteria for untrained singers," in *2008 9th International Conference on Signal Processing*. IEEE, 2008, pp. 1475–1478.
- [5] Chitrlekha Gupta, Haizhou Li, and Ye Wang, "A technical framework for automatic perceptual evaluation of singing quality," *APSIPA Transactions on Signal and Information Processing*, vol. 7, 2018.
- [6] Chitrlekha Gupta, Haizhou Li, and Ye Wang, "Perceptual evaluation of singing quality," in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2017, pp. 577–586.
- [7] Koichi Omori, Ashutosh Kacker, Linda M Carroll, William D Riley, and Stanley M Blaugrund, "Singing power ratio: quantitative evaluation of singing voice quality," *Journal of Voice*, vol. 10, no. 3, pp. 228–235, 1996.
- [8] Chuan Cao, Ming Li, Jian Liu, and Yonghong Yan, "An objective singing evaluation approach by relating acoustic measurements to perceptual ratings," in *Proceedings of INTERSPEECH 2008*, 2008, pp. 2058–2061.
- [9] Yang Yu, Weisi Lin, Dong-Yan Huang, Minghui Dong, and Haizhou Li, "Performance scoring of singing voice," in *2015 International Conference on Asian Language Processing (IALP)*. IEEE, 2015, pp. 119–122.
- [10] Edward Pórolniczak and Michał Kramarczyk, "Computer assessment of tremolo feature in the context of evaluation of singing quality," in *2015 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*. IEEE, 2015, pp. 157–161.
- [11] Terry Tan, "Singing evaluation based on deep metric learning," in *Proceedings of the 2019 3rd International Symposium on Computer Science and Intelligent Control*, 2019, pp. 1–5.
- [12] Huan Zhang, Yiliang Jiang, Tao Jiang, and Hu Peng, "Learn by referencing: Towards deep metric learning for singing assessment," in *the 22nd International Society for Music Information Retrieval Conference (ISMIR)*, 2021, pp. 825–832.
- [13] Tomoyasu Nakano, Masataka Goto, and Yuzuru Hiraga, "An automatic singing skill evaluation method for unknown melodies using pitch interval accuracy and vibrato features," in *Proceedings of INTERSPEECH 2006*, 2006, pp. 1706–1709.
- [14] Chitrlekha Gupta, Haizhou Li, and Ye Wang, "Automatic leaderboard: Evaluation of singing quality without a standard reference," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 13–26, 2019.
- [15] Ning Zhang, Tao Jiang, Feng Deng, and Yan Li, "Automatic singing evaluation without reference melody using bi-dense neural network," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 466–470.
- [16] Chitrlekha Gupta, Lin Huang, and Haizhou Li, "Automatic rank-ordering of singing vocals with twin-neural network," in *the 21st International Society for Music Information Retrieval Conference (ISMIR)*, 2020, pp. 416–423.
- [17] Lin Huang, Chitrlekha Gupta, and Haizhou Li, "Spectral features and pitch histogram for automatic singing quality evaluation with crnn," in *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2020, pp. 492–499.
- [18] Jinhu Li, Chitrlekha Gupta, and Haizhou Li, "Training explainable singing quality assessment network with augmented data," in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2021, pp. 904–911.
- [19] Chitrlekha Gupta, Jinhu Li, and Haizhou Li, "Towards reference-independent rhythm assessment of solo singing," in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2021, pp. 912–919.
- [20] Joel Wapnick and Elizabeth Ekholm, "Expert consensus in solo voice performance evaluation," *Journal of Voice*, vol. 11, no. 4, pp. 429–436, 1997.
- [21] Stephen McAdams and Bruno L Giordano, "The perception of musical timbre," *The Oxford handbook of music psychology*, pp. 72–80, 2009.
- [22] Müller Meinard and Frank Zalkow, "Fmp notebooks: Educational material for teaching and learning fundamentals of music processing," in *the 20th International Society for Music Information Retrieval Conference (ISMIR)*, 2019, pp. 573–580.
- [23] Kyungyun Lee and Juhan Nam, "Learning a joint embedding space of monophonic and mixed music signals for singing voice," in *the 20th International Society for Music Information Retrieval Conference (ISMIR)*, 2019, pp. 295–302.
- [24] JMH van Balen, John Ashley Burgoyne, Frans Wiering, and Remco C Veltkamp, "An analysis of chorus features in popular song," in *the 14th International Society of Music Information Retrieval Conference (ISMIR)*, 2013, pp. 107–112.
- [25] Laurens Van der Maaten and Geoffrey Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [26] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, et al., "Deep high-resolution representation learning for visual recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 10, pp. 3349–3364, 2020.
- [27] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter, "Fast and accurate deep network learning by exponential linear units (elus)," *arXiv preprint arXiv:1511.07289*, 2015.
- [28] Zhiyan Duan, Haotian Fang, Bo Li, Khe Chai Sim, and Ye Wang, "The nus sung and spoken lyrics corpus: A quantitative comparison of singing and speech," in *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2013, pp. 1–9.