

MULTIPLE F0 ESTIMATION IN VOCAL ENSEMBLES USING CONVOLUTIONAL NEURAL NETWORKS

Helena Cuesta¹ Brian McFee² Emilia Gómez^{1,3}

¹ Music Technology Group, Universitat Pompeu Fabra, Barcelona

² Music and Audio Research Lab & Center for Data Science, New York University, USA

³ European Commission, Joint Research Centre, Seville

helena.cuesta@upf.edu

ABSTRACT

This paper addresses the extraction of multiple F0 values from polyphonic and *a cappella* vocal performances using convolutional neural networks (CNNs). We address the major challenges of ensemble singing, i.e., all melodic sources are vocals and singers sing in harmony. We build upon an existing architecture to produce a pitch salience function of the input signal, where the harmonic constant-Q transform (HCQT) and its associated phase differentials are used as an input representation. The pitch salience function is subsequently thresholded to obtain a multiple F0 estimation output. For training, we build a dataset that comprises several multi-track datasets of vocal quartets with F0 annotations. This work proposes and evaluates a set of CNNs for this task in diverse scenarios and data configurations, including recordings with additional reverb. Our models outperform a state-of-the-art method intended for the same music genre when evaluated with an increased F0 resolution, as well as a general-purpose method for multi-F0 estimation. We conclude with a discussion on future research directions.

1. INTRODUCTION

Ensemble singing is a well-established practice across cultures, found in a great diversity of forms, languages, and levels. However, all variants share the social aspect of collective singing, either as a form of entertainment or expressing emotions. In *The Science of the Singing Voice* [1], Sundberg claims that choral singing is one of the most widespread types of singing. In Western classical music, a *choir* is usually a group of singers divided into four sections: soprano, alto, tenor, bass (SATB); however, there exist many other forms of polyphonic singing, involving a diverse number of singers, parts, and vocal ranges. One example of such variants is a vocal quartet, where four singers—commonly with distinct vocal ranges—sing in

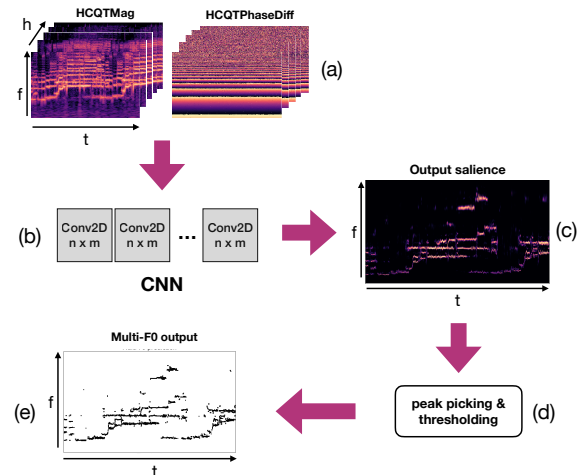


Figure 1. Overview of the proposed method. (a) Input features. (b) Convolutional architecture diagram. (c) Output salience map. (d) Peak picking and thresholding step to obtain F0 values from the salience activation in (c). (e) Multiple F0 representation, output of the framework.

harmony. Vocal quartets usually follow the SATB configuration; therefore, they are different from a *standard* choir in that there is only one singer per section.

Ensemble singing has not been widely studied in the field of Music Information Retrieval (MIR) in the recent years. We find a few early works focused on the acoustic properties of choral singing [2–4], and also a few more recent studies about some expressive characteristics of polyphonic vocal music, such as singer interaction and intonation [5–10], analysis of unison singing [11, 12], or choir source separation [13]. Most of these studies rely on individual recordings of each voice in the ensemble, which enable the automatic extraction of fundamental frequency (F0) contours from each isolated voice. The general applicability of these approaches is limited by the fact that vocal groups are rarely recorded with individual microphones. Another common strategy is using the polyphonic audio recordings along with their associated synchronized scores. However, the process of synchronizing a choral audio recording to a score is not straightforward, and therefore results are not always entirely trustworthy. Manual an-



© H. Cuesta, B. McFee, and E. Gómez. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

Attribution: H. Cuesta, B. McFee, and E. Gómez, “Multiple F0 Estimation in Vocal Ensembles using Convolutional Neural Networks”, in *Proc. of the 21st Int. Society for Music Information Retrieval Conf.*, Montréal, Canada, 2020.

notation is also one solution to obtain reliable F0 contours, but the process is highly time-consuming and expensive.

An alternate approach is to analyze a mixed, polyphonic recording to produce multiple F0 estimations simultaneously. This process enables the use of polyphonic singing recordings in the wild, and eliminates the need for separate audio recordings for each singer. However, multi-F0 estimation in polyphonic vocal music has been less often studied, likely due to the complexity and variety of sounds that a singer can produce, the timbre similarity between singers' voices, as well as the scarcity of annotated data of this kind [14].

In this paper, we focus on multiple F0 estimation for vocal ensemble audio mixtures. We build upon previous work on using neural networks to obtain an intermediate salience representation suitable for several tasks, including multi-F0 estimation [15]. In particular, we experiment with a set of convolutional neural networks (CNN), and combine magnitude and phase information, which is commonly neglected in the literature. We experiment with different fusion strategies and analyze the generalization capabilities of deep learning models in the presence of unison and reverbs. Figure 1 shows a diagram of the proposed method.

Following research reproducibility principles, data generation scripts and models are accessible¹.

2. RELATED WORK

Multiple F0 estimation (also referred to as multi-F0 estimation, or multi-pitch estimation) is a sub-task of automatic music transcription (ATM) that consists of detecting multiple concurrent F0 values in an audio signal that contains several melodic lines at the same time [14, 16, 17]. Benetos et al. [17] summarize the main challenges of ATM as follows: polyphonic mixtures having multiple simultaneous sources with different pitches, loudness, and timbre properties; sources with overlapping harmonics; and the lack of polyphonic music datasets with reliable ground truth annotations, among others. They organize ATM approaches into four categories: frame-level (or multi-F0 estimation), note-level (or note-tracking), stream-level (or multi-F0 streaming), and notation-level. Our work focuses on the first category.

While monophonic F0 estimation is a well-researched topic, with state-of-the-art systems with excellent performances [18–21], multiple F0 estimation is still challenging. Research on this topic is commonly divided into several groups according to the nature of the employed methods. For instance, in [17] they report four categories: traditional signal processing methods, probabilistic methods, non-negative matrix factorization (NMF) methods, and neural networks.

Klapuri [22] proposed a signal processing based method for multi-F0 estimation in polyphonic music. He calculates the salience of F0 candidates by summing the amplitudes of its harmonic partials, and then uses an iterative method

where at every step an F0 is estimated and cancelled from the mixture before moving to the next iteration to estimate the next F0. The same author presented in [23] a similar method that incorporates information about human perception by means of an auditory model before the iterative process.

The system presented by Duan et al. [24] uses maximum-likelihood approach with the power spectrum as input. Spectral peaks are detected and two separate regions are defined accordingly: the peak region and the non-peak region, using a tolerance of half semitone from the detected peaks. In the maximum-likelihood process, both sets are treated independently, and the process of detecting F0 consists of optimizing a joint function that maximizes the probability of having harmonics that explain the observed peaks and minimizing the probability of having harmonics in the non-peak region. The F0 estimates are post-processed using neighbouring frames' estimates to produce more stable F0 contours.

A recent example of a multiple F0 estimation framework that employs neural networks is the system by Bittner et al. [15], DeepSalience (DS): a CNN trained to produce a multi-purpose pitch salience representation of the input signal. It is designed for multi-instrument pop/rock polyphonic music, and it provides an intermediate representation for MIR tasks such as melody extraction and multi-F0 estimation, outperforming state-of-the-art approaches in both cases. Following the premise that a pitch salience function is a suitable representation to extract F0 values, also exploited in [22, 23, 25, 26], this work keeps up with the advancements of deep neural networks to build data-driven salience functions. They use the harmonic constant-Q transform (HCQT) as input feature, which is a 3-dimensional array indexed by harmonic index h , frequency f , and time t : $\mathcal{H}[h, f, t]$. It comprises a set of constant-Q transforms (CQT) stacked together, each of them with its minimum frequency scaled by the harmonic index: $h \cdot f_{min}$.

While the above methods are well-suited for multiple F0 estimation in multi-instrumental music, *a cappella* polyphonic vocal music has several particularities that justify the need for dedicated techniques. One of the most significant challenges of analyzing vocal ensembles is due to harmonies occurring between distinct, overlapping vocal ranges. The timbre similarity, strong harmonic relationships, and overlapping frequency ranges hinder the extraction of concurrent F0 values in such music signals.

McLeod et al. [27] present a system for automatic transcription of polyphonic vocal music, which includes an initial step of estimating multiple F0s, and a second step of voice assignment, where each detected F0 is assigned to one of the SATB voices. They combine an acoustic model based on the factorization of an input log-frequency spectrogram for the multi-F0 estimation with a music language model based on hidden Markov models (HMM) for the voice assignment step. An earlier version of this method was presented in [14]; however, in the latter work the authors include a model integration step where the output of

¹Companion code and models: <https://github.com/helenacuesta/multif0-estimation-vocals>

Dataset	# of songs	Duration (hh:mm:ss)
Choral Singing Dataset [8]	3 songs	00:07:14
Dagstuhl ChoirSet [29]	2 songs	00:55:30
ESMUC Choir Dataset	3 songs	00:21:08
Barbershop Quartets ²	22 songs	00:42:10
Bach Chorales ³	26 songs	00:58:20

Table 1: Overview of the datasets used in this paper. The reported durations refer to the original mixtures before re-mixing stems and data augmentation. Dagstuhl ChoirSet and ESMUC Choir Dataset contain several takes per song.

the music language model is further used in the acoustic model to improve the estimation of F0s. Their results show that integrating both parts of the system improves the performance of the voice assignment, and also of the multi-F0 estimation, since it eliminates many false positives.

Su et al. [28] also address some of the aforementioned issues by proposing an unsupervised method for multi-F0 estimation of choir and symphonic music. Their approach uses time-frequency reassignment techniques such as the synchrosqueezing transform (SST), which aims to better discriminate closely-located spectral components, such as unisons. They use an improved technique called ConceFT, which is based on the idea of multi-taper SST, but was proved to estimate instantaneous frequencies in noisy signals more precisely. These methods measure pitch salience and enhance the stability and localization of the F0 features needed for multi-F0 estimation.

3. DATASET

The lack of an appropriate and large enough annotated dataset has been a bottleneck in the use of machine learning techniques for multiple F0 estimation in ensemble singing. We address this difficulty by constructing a dataset that comprises several multi-track datasets of polyphonic singing with F0 annotations. We created a dataset by aggregating several existing multi-track polyphonic singing datasets. Table 1 shows an overview of the characteristics of each dataset individually. In this section, we describe them in more detail, as well as explain the process of data augmentation.

We use five datasets of similar characteristics. First, the Choral Singing Dataset (CSD) [8], a publicly available multi-track dataset of Western choral music. It comprises recordings of three SATB songs performed by a choir of 16 singers, four per section (4S4A4T4B), and it contains separate audio stems for each singer. Besides, it includes F0 annotations for each singer, which are automatically extracted and manually corrected. Similarly, the ESMUC Choir Dataset (ECS) is a proprietary dataset that comprises three songs performed by a choir of 13 singers (5S3A3T2B); it also includes audio stems for each

² <https://www.pgmusic.com/barbershopquartet.htm>

³ <https://www.pgmusic.com/bachchorales.htm>

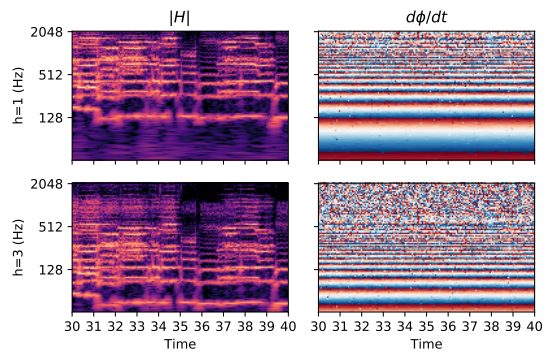


Figure 2. Input features examples for 10-second excerpt of a file from the training set. HCQT magnitude is depicted in the left column, and HCQT phase differentials in the right column, for $h = 1$ (top) and $h = 3$ (bottom).

singer and F0 annotations. The third dataset with comparable characteristics is the Dagstuhl ChoirSet (DCS) [29]: it consists of recordings of two songs performed by a choir of 13 singers (2S2A4T5B), and two different SATB quartets. This dataset also provides the audio stems and automatically extracted F0 annotations. Finally, we also add two commercial datasets: the Bach Chorales (BC)³ and the Barbershop quartets (BSQ)². They contain 26 and 22 songs, respectively, performed by vocal quartets—SATB in the first case, and tenor, lead, baritone, bass in the second case—as well as automatically extracted F0 annotations.

We exploit the multi-track nature of all datasets to create artificial mixtures of stems. We use PySox [30] to create all the possible combinations of singers, with the constraint of having one singer per part (SATB). In parallel, we also generate the multi-F0 annotations by combining the individual F0 contours of each singer in the mixture.

Besides creating the audio mixtures from individual recordings, we include two additional steps to improve generalization. First, we augment our dataset by means of pitch-shifting individual voices and re-mixing them. Particularly, we use pitch-shifting at a semitone scale: -2 to $+2$ semitones from the original signal. Second, our dataset contains two versions of each audio clip: the original one (obtained by mixing together individual stems), and the same song with reverb. We use the *Great Hall* impulse response (IR) from the Room Impulse Response Dataset in Isophonics [31], and convolve it with the audio mixtures of our dataset. For both tasks, we use MUDA, a software framework for musical data augmentation [32].

The dataset consists of 22910 audio files of diverse durations, from 10 seconds to 3 minutes. We split it into training (75%, 17184 files), validation (10%, 2291 files), and test (15%, 3435 files) subsets.

4. PROPOSED METHOD

In this section, we describe the input features, the target representations, the convolutional architectures we design, and the experiments we conduct.

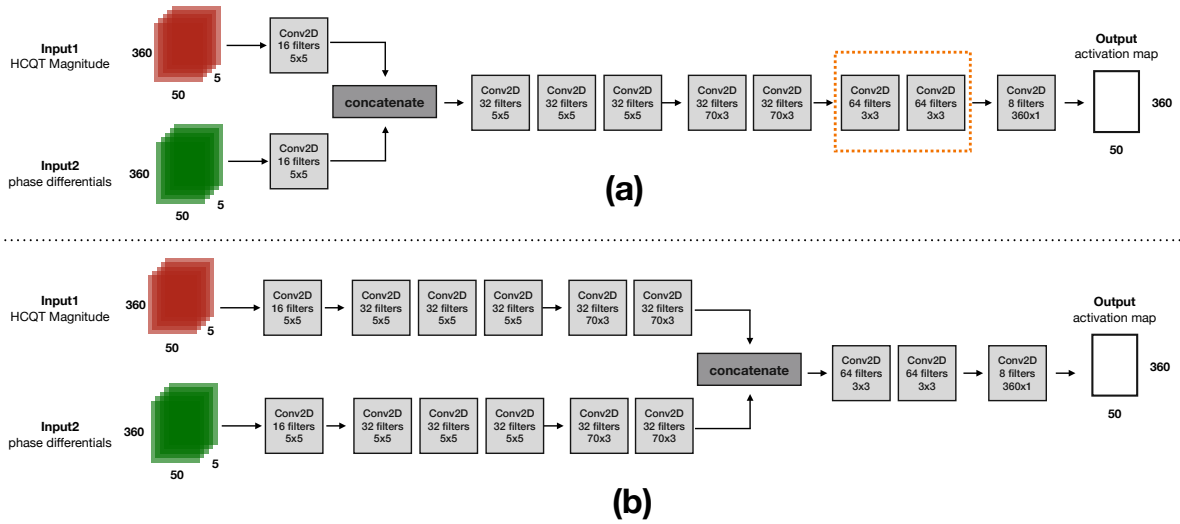


Figure 3. Proposed convolutional architectures. **(a)** Early/Shallow and Early/Deep: the two layers inside the orange dotted rectangle are only part of Early/Deep. **(b)** Late/Deep: the concatenation of both inputs’ contribution happens later in the network. In all networks, each layer is preceded by a batch-normalization step and the output of each layer is passed through a rectified linear unit activation function, except for the last layer, which uses a sigmoid.

4.1 Input features

Our networks have two separate inputs: the HCQT magnitude and the HCQT phase differentials. The HCQT is a 3-dimensional array $\mathcal{H}[h, t, f]$, indexed by harmonic (h), frequency (f), and time (t). It measures the h th harmonic of frequency f at time t , where $h = 1$ is the fundamental. This representation is based on computing a standard constant-Q transform (CQT) for each harmonic where the minimum frequency (f_{min}) is scaled by the harmonic number, $h \cdot f_{min}$. Detailed descriptions of the HCQT are presented in [15, 33]. For the HCQT calculation we use 60 bins per octave, 20 cents per bin, 6 octaves, and a minimum frequency of 32.7 Hz, which corresponds to a C1. We use five harmonics, so that $h \in \{1, 2, 3, 4, 5\}$ to compute the frequencies of the partials. Phase information is often discarded from neural network inputs, which commonly use magnitude representations such as the magnitude of the short-time Fourier transform (STFT). However, we also use the associated phase differentials. From signal processing theory we know that the phase differential of a signal contributes to a more precise calculation of the instantaneous frequency (ω_{ins}) [34]:

$$\omega_{ins} = \frac{\delta\varphi(t)}{\delta t} \rightarrow f_{ins} = \frac{1}{2\pi} \frac{\delta\varphi(t)}{\delta t} \quad (1)$$

where $\varphi(t)$ is the phase spectrum of the audio signal.

All audio files are resampled to a sampling rate of 22050 Hz, and we use a hop size of 256 samples. An example of the two input features examples is displayed in Figure 2.

4.2 Output representation

The output targets we use to train our networks are time-frequency representations with the same dimensions (2-D) as one of the input channels, i.e., $\mathcal{H}[1]$. We use the

ground truth F0 annotations (see Section 3) and assign each F0 value to the nearest time-frequency bin in the 2-D representation—which has the same time and frequency resolutions as the input—with a magnitude of 1. Non-active bins are set to 0, and we apply Gaussian blur with standard deviation 1 in the frequency direction to account for possible imprecisions in the predictions. We follow the same procedure as in DeepSaliency and set the energy decay from 1 to 0 to cover half a semitone in frequency.

4.3 Models

Figure 3 depicts three convolutional architectures we propose: Early/Shallow, Early/Deep, and Late/Deep.

4.3.1 Early/Shallow and Early/Deep models

These models, inspired by DeepSaliency, are illustrated in Figure 3a. They both consist of a fully convolutional architecture with two separate inputs: one for the HCQT magnitude and a second one for the HCQT phase differentials. Each of these inputs is first sent to a convolutional layer with 16 (5×5) filters. Then, the outputs of these two layers are concatenated. (5×5) filters cover approximately 1 semitone in frequency and 50 ms in time. After the concatenation, data passes through a set of convolutional layers including two layers with 32 (70×3) filters, which cover 14 semitones in frequency and are suitable for capturing harmonic relations within an octave. In the Early/Deep model we add two 64 (3×3) layers before the last layer with 8 filters that cover all frequency bins.

4.3.2 Late/Deep model

Late/Deep diagram is displayed in Figure 3b, and it follows a similar structure to Early/Shallow and Early/Deep. However, in this case both inputs are handled separately

until the layer with (70×3) filters; then, we concatenate both data streams and add the same layers: two layers with 64 filters (3×3) , and the last layer with 8 filters that cover the whole frequency dimension, i.e., 360 bins.

In all models, batch normalization is applied at the input of every layer, and the outputs are passed through rectified linear units (ReLU), except for the output layer, which uses logistic activation (sigmoid) to map the output of each bin to the range $[0, 1]$. Using sigmoid at the output enables the interpretation of the activation map as a probability function, where the value between 0 and 1 represents the probability that a specific bin belongs to the set of F0s present in the input signal. All models in the experiments described next are trained to minimize binary cross-entropy between the target, $y[t, f]$, and the prediction, $\hat{y}[t, f]$, both of them values in the range $[0, 1]$:

$$L(y, \hat{y}) = -y \log(\hat{y}) - (1 - y) \log(1 - \hat{y}) \quad (2)$$

We use the Adam optimizer [35] with a learning rate of 0.001, and train for 100 epochs with a batch size of 16 patches of shape $(360, 50)$. We perform early stopping when the validation error does not decrease for 25 epochs.

4.4 Experimental setup

4.4.1 Evaluation metrics

We evaluate the models using the frame-wise metrics Precision, Recall, F-Score (or F-measure) as they are defined in the MIREX multiple-F0 estimation task [36] using the `mir_eval` library [37].

4.4.2 Experiment 1: fusion strategy

In the first experiment we use the whole dataset split into train-validation-test subsets (see Section 3), to measure the general performance of the three models. With this experiment we study the influence of magnitude and phase information fusion at an early stage of the network, i.e., Early/Shallow-Deep, or later, i.e., Late/Deep. For each model, we use the validation set to optimize the threshold we apply to the peaks extracted from the output salience representation. The optimal threshold is the one that maximizes the average accuracy across the validation set in each case. In addition, we train the Late/Deep model without the phase information, i.e., we remove the branch of the network dedicated to the phase. We intend to verify the hypothesis that including the phase as input to the network leads to more precise results.

4.4.3 Experiment 2: comparative analysis

In this experiment we evaluate the performance of our best-performing model from Experiment 1 on the BSQ². This is one of the datasets used in [27, 38], allowing for a direct comparison between their method—also designed for ensemble singing—and the model we propose. These data are part of our original training dataset, but in this experiment we train the model excluding all the BSQ audio files, and then use them for exclusively for evaluation.

4.4.4 Experiment 3: generalization

In this last experiment, we aim to explore the effect of unison and reverb. Since vocal ensembles are commonly captured using a room microphone, such recordings usually contain reverb or similar effects, caused by the room acoustics. We train our best-performing model excluding all audio files with reverb from the dataset, and then evaluate it with conventional choir recordings from the dataset presented in [28], which is not part of our working dataset. In addition, we evaluate this model on a subset of reverb files from the original test set and compare the performance of this model to the model trained in Experiment 1.

5. RESULTS

5.1 Experiment 1: fusion strategy

Results for Experiment 1 are depicted in Figure 4. While the three models have similar results, Late/Deep is slightly better in terms of F-Score, suggesting that the late fusion of magnitude and phase information is more robust than the early fusion. Figure 5 shows an excerpt of the multiple F0 output (red) together with their associated ground truth reference (black). We compare these results to DeepSalience using a detection threshold of 0.2 (optimized beforehand on the evaluation material), which has a lower performance, which we attribute to distribution shift from its training set. Additionally, we observe how the Late/Deep without phase information has a similar F-Score but lower precision, showing that including phase differentials as input is helpful for obtaining more accurate results. Since Late/Deep is the model with the best performance, we use it in the subsequent experiments.

5.2 Experiment 2: comparative analysis

Experiment 2 results are summarized in Table 2 in terms of F-Score, and Precision and Recall (when available). We evaluate the predictions from our model on the BSQ dataset, and compare the results to the ones reported with MSINGERS [14], and VOCAL4-VA, the fully-integrated model from [27]. We use two different pitch tolerances: one semitone (100 cents) and 20 cents. While one semitone resolution is enough for transcription purposes, for analysis such as the ones described in Section 1, i.e., intonation and singer interaction, more pitch resolution is required. We observe that our model outperforms both baseline methods with two different pitch tolerances. In the 20 cents evaluation, the baseline models experience a performance drop (-17% and -26%), whereas the decrease in our model is much smaller, around -2% in the F-Score. This difference presumably resides in the fact that our model uses phase information to refine F0 estimates, therefore extracting a more precise value.

5.3 Experiment 3: generalization

Results from this experiment show that our model outperforms the method in [28] on their choir dataset: when we




Figure 4. Evaluation results on the non-pitch-shifted audio files of the test set. We compare our three models to DeepSaliency as a baseline, as well as to the Late/Deep network trained without the phase differentials (Late/Deep no-phase). Note that outliers are excluded from the plots for an easier visualization.




Figure 5. Example output from Experiment 1, Late/Deep model. Predictions (red) are plotted over the ground truth reference (black).

Method	100 cents			20 cents		
	F	P	R	F	P	R
MSINGERS [14]	0.708 (0.06)	0.685 (0.06)	0.736 (0.07)	0.537 (0.07)	0.620 (0.07)	0.477 (0.08)
VOCAL4-VA [27]	0.757 (0.06)	-	-	0.490	-	-
Late/Deep	0.846 (0.03)	0.812 (0.03)	0.884 (0.04)	0.831 (0.03)	0.797 (0.03)	0.868 (0.04)

Table 2: Multi-F0 estimation results (F-Score (F), precision (P), and recall (R)) on the Barbershop quartets, for different pitch tolerances. Values in parentheses refer to the standard deviation. Best scores are highlighted in bold.

calculate the average F-Score across the whole dataset, using the threshold optimized on the validation set, we obtain 0.704, while their best-performing method reaches an average F-Score of 0.653. Note that this dataset contains short excerpts of commercial choir recordings, with several singers per section and a large reverb effect, which differs from our training material. Therefore, these results suggest that our model is robust to recordings in such context. However, a larger experiment with similar data would be necessary, since this dataset is very small. The second part of this experiment is the evaluation of a subset of ten reverb files from the test set. In terms of F-Score, and as expected, the model that includes reverb files in the training set (Experiment 1) improves by slightly more than 10% on

average with respect to the model that excludes these files (Experiment 3). Therefore, we conclude that the presence of both reverb and dry signals in the training set is beneficial for the performance of a wider range of recording conditions such as reverb.

6. CONCLUSIONS

In this paper, we proposed a set of novel convolutional architectures for multiple F0 estimation in *a cappella* ensemble singing, combining magnitude and phase information. For training, we created an annotated dataset of polyphonic singing voice by aggregating several existing datasets, and augmented it by means of pitch-shifting and reverberation.

We conducted several experiments to evaluate different aspects of the detection process. We evaluated the overall performance of three models as compared to a deep learning based multi-purpose multiple F0 estimation system, and found that our models outperform the baseline when applied on ensemble singing. We also verified that using phase information at the input, together with the magnitude, improves the precision of the F0 estimates. We compared our best-performing model to one existing approach for multiple F0 estimation in vocal ensembles, and demonstrated that it outperforms it with two different F0 resolutions (100 and 20 cents). In addition, we compared our model to an approach specifically designed for choir and symphonic music and found that our model is robust in conditions of unison and high reverb. However, further experiments with a larger amount of data are required to verify these findings.

Although our results are a strong contribution to addressing the limitations of deep learning architectures for vocal music, there are some further steps that would potentially improve the performance of our models. Informal experiments showed that post-processing the output F0 contours increases their time continuity; therefore, the overall quality of the output improves. Further steps also include not only estimating the F0 values frame-wise, but also assigning each of them to a singer, which is a challenging task if the number of singers is not known a priori.

7. ACKNOWLEDGEMENTS

The authors would like to thank Rodrigo Schramm and Emmanouil Benetos for sharing the BSQ and BC datasets for this research. Helena Cuesta is supported by the FI Pre-doctoral Grant from AGAUR (Generalitat de Catalunya). This work is partially supported by the European Commission under the TROMPA project (H2020 770376) and MARL-NYU (as part of a two-months research stay).

8. REFERENCES

- [1] J. Sundberg, *The Science of the Singing Voice*. Northern Illinois University Press, 1987.
- [2] T. D. Rossing, J. Sundberg, and S. Ternström, “Acoustic comparison of voice use in solo and choir singing,” *The Journal of the Acoustical Society of America*, vol. 79, no. 6, pp. 1975–1981, 1986.
- [3] S. Ternström, “Perceptual evaluations of voice scatter in unison choir sounds,” *STL-Quarterly Progress and Status Report*, vol. 32, pp. 041–049, 1991.
- [4] S. Ternström, “Choir acoustics – an overview of scientific research published to date,” *Speech, Music and Hearing Quarterly Progress and Status Report*, vol. 43, no. April, pp. 001–008, 2002.
- [5] S. Zadig, G. Folkestad, and V. Lyberg-Ahlander, “Multi-track recordings of choral singers: Development and validation of a method to identify activities and interaction in the choral voice, based on recordings of the individual singers,” *Bulletin of Empirical Music Education Research*, vol. 7, no. 1, pp. 1–20, 2016.
- [6] J. Devaney, M. I. Mandel, and I. Fujinaga, “A study of intonation in three-part singing using the automatic music performance analysis and comparison toolkit (AMPACT),” in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, Porto, Portugal, 2012, pp. 511–516.
- [7] J. Dai and S. Dixon, “Analysis of interactive intonation in unaccompanied SATB ensembles,” in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, Suzhou, China, 2017, pp. 599–605.
- [8] H. Cuesta, E. Gómez, A. Martorell, and F. Loáiciga, “Analysis of intonation in unison choir singing,” in *Proceedings of the International Conference of Music Perception and Cognition (ICMPC)*, Graz, Austria, 2018, pp. 125–130.
- [9] J. Dai and S. Dixon, “Singing together: Pitch accuracy and interaction in unaccompanied unison and duet singing,” *The Journal of the Acoustical Society of America*, vol. 145, no. 2, pp. 663–675, 2019.
- [10] C. Weiss, S. J. Schelcht, S. Rosenzweig, and M. Müller, “Towards Measuring Intonation Quality of Choir Recordings: A Case Study on Bruckner’s Locus Iste,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR)*, Delft, The Netherlands, 2019, pp. 276–283.
- [11] H. Cuesta, E. Gómez, and P. Chandna, “A framework for multi- f_0 modeling in satb choir recordings,” in *Proceedings of the Sound and Music Computing Conference (SMC)*, Málaga, Spain, 2019, pp. 447–453.
- [12] P. Chandna, H. Cuesta, and E. Gómez, “A deep learning based analysis-synthesis framework for unison singing,” in *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, 2020.
- [13] D. Petermann, P. Chandna, H. Cuesta, J. Bonada, and E. Gómez, “Deep learning based source separation applied to choir ensembles,” in *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, 2020.
- [14] R. Schramm and E. Benetos, “Automatic transcription of a cappella recordings from multiple singers,” in *Proceedings of the AES Conference on Semantic Audio*, Erlangen, Germany, 2017.
- [15] R. M. Bittner, B. McFee, J. Salamon, P. Li, and J. P. Bello, “Deep salience representations for F0 tracking in polyphonic music,” in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, Suzhou, China, 2017, pp. 63–70.
- [16] E. Benetos, S. Dixon, D. Giannoulis, H. Kirchhoff, and A. Klapuri, “Automatic music transcription: challenges and future directions,” *Journal of Intelligent Information Systems*, vol. 41, no. 3, pp. 407–434, 2013.
- [17] E. Benetos, S. Dixon, and Z. Duan, “Automatic Music Transcription : An Overview,” *IEEE Signal Processing Magazine*, vol. 36, no. 1, pp. 20–30, 2018.
- [18] A. D. Cheveigné and H. Kawahara, “Yin, a fundamental frequency estimator for speech and music,” *The Journal of the Acoustical Society of America*, vol. 111, no. 4, pp. 1917–1930, 2002.
- [19] A. Camacho and J. G. Harris, “A sawtooth waveform inspired pitch estimator for speech and music,” *The Journal of the Acoustical Society of America*, vol. 124, no. 3, pp. 1638–1652, 2008.
- [20] M. Mauch and S. Dixon, “pYIN: A fundamental frequency estimator using probabilistic threshold distributions,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Florence, Italy, 2014, pp. 659–663.
- [21] J. W. Kim, J. Salamon, P. Li, and J. P. Bello, “CREPE: A Convolutional Representation for Pitch Estimation,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, 2018, pp. 161–165.

- [22] A. P. Klapuri, "Multiple fundamental frequency estimation by summing harmonic amplitudes," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2006, pp. 216–221.
- [23] A. Klapuri, "Multipitch analysis of polyphonic music and speech signals using an auditory model," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 2, pp. 255–266, Feb 2008.
- [24] Z. Duan, B. Pardo, and C. Zhang, "Multiple Fundamental Frequency Estimation by Modeling Spectral Peaks and Non-peak Regions," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 8, pp. 2121–2133, 2010.
- [25] M. P. Rynnänen and A. P. Klapuri, "Automatic transcription of melody, bass line, and chords in polyphonic music," *Computer Music Journal*, vol. 32, no. 3, pp. 72–86, 2008.
- [26] J. Salamon and E. Gómez, "Melody extraction from polyphonic music signals using pitch contour characteristics," *IEEE Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 20, pp. 1759–1770, 08 2012.
- [27] A. McLeod, R. Schramm, M. Steedman, and E. Benetos, "Automatic transcription of polyphonic vocal music," *Applied Sciences*, vol. 7, no. 12, 2017.
- [28] L. Su, T.-Y. Chuang, and Y.-H. Yang, "Exploiting frequency, periodicity and harmonicity using advanced time-frequency concentration techniques for multipitch estimation of choir and symphony," in *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, New York City, USA, 2016, pp. 393–399.
- [29] S. Rosenzweig, H. Cuesta, C. Weiss, F. Scherbaum, E. Gómez, and M. Müller, "Dagstuhl ChoirSet: A Multitrack Dataset for MIR Research on Choral Singing," *Transactions of the International Society for Music Information Retrieval (TISMIR)*, vol. 3, no. 1, pp. 98–110, 2020.
- [30] R. M. Bittner, E. Humphrey, and J. P. Bello, "Pysox: Leveraging the audio signal processing power of sox in python," in *Proceedings of the International Society for Music Information Retrieval Conference Late Breaking and Demo Papers*, 2016.
- [31] R. Stewart and M. Sandler, "Database of omnidirectional and b-format room impulse responses," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2010, pp. 165–168.
- [32] B. McFee, E. J. Humphrey, and J. P. Bello, "A software framework for musical data augmentation," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2015, pp. 248–254.
- [33] R. M. Bittner, B. McFee, and J. P. Bello, "Multitask learning for fundamental frequency estimation in music," *ArXiv*, vol. abs/1809.00381, 2018.
- [34] B. Boashash, "Estimating and Interpreting the Instantaneous Frequency of a Signal. I. Fundamentals," *Proceedings of the IEEE*, vol. 80, no. 4, pp. 520–538, 1992.
- [35] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *ArXiv*, vol. abs/1412.6980, 2014.
- [36] M. Bay, A. F. Ehmann, and S. J. Downie, "Evaluation of multiple-f0 estimation and tracking systems," in *Proceedings of the 10th International Society for Music Information Retrieval Conference (ISMIR)*, 2009, pp. 315–320.
- [37] C. Raffel, B. McFee, E. J. Humphrey, J. Salamon, O. Nieto, D. Liang, and D. P. Ellis, "mir_eval: A transparent implementation of common mir metrics," in *In Proceedings of the 15th International Society for Music Information Retrieval Conference (ISMIR)*, 2014.
- [38] R. Schramm, A. McLeod, M. Steedman, and E. Benetos, "Multi-pitch detection and voice assignment for a cappella recordings of multiple singers," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2017.