

A Framework for Multi- f_0 Modeling in SATB Choir Recordings

Helena Cuesta

Music Technology Group
Universitat Pompeu Fabra
helena.cuesta@upf.edu

Emilia Gómez

Universitat Pompeu Fabra and
Joint Research Centre (EC)
emilia.gomez@upf.edu

Pritish Chandna

Music Technology Group
Universitat Pompeu Fabra
pritish.chandna@upf.edu

ABSTRACT

Fundamental frequency (f_0) modeling is an important but relatively unexplored aspect of choir singing. Performance evaluation as well as auditory analysis of singing, whether individually or in a choir, often depend on extracting f_0 contours for the singing voice. However, due to the large number of singers, singing at a similar frequency range, extracting the exact individual pitch contours from choir recordings is a challenging task. In this paper, we address this task and develop a methodology for modeling pitch contours of SATB choir recordings. A typical SATB choir consists of four parts, each covering a distinct range of pitches and often with multiple singers each. We first evaluate some state-of-the-art multi- f_0 estimation systems for the particular case of choirs with a single singer per part, and observe that the pitch of individual singers can be estimated to a relatively high degree of accuracy. We observe, however, that the scenario of multiple singers for each choir part (i.e. unison singing) is far more challenging. In this work we propose a methodology based on combining a multi- f_0 estimation methodology based on deep learning followed by a set of traditional DSP techniques to model f_0 and its dispersion instead of a single f_0 trajectory for each choir part. We present and discuss our observations and test our framework with different singer configurations.

1. INTRODUCTION

Singing in a SATB (Soprano, Alto, Tenor, Bass) choir is a long standing and well enjoyed practice, with many choirs following this format across different languages and cultures. Performances are based on scores, which provide linguistic, timing and pitch information for the singers in the choir to follow. Professional choirs practice for years to ensure that their performance is *in tune* with a reference pitch; however, due to the mechanism of voice production and expressive characteristics, the pitch of the individual voices in the choir often deviates from the theoretical pitch as indicated in the score. As a consequence, analysis and evaluation of a choir performance depends on the combination of pitches produced by individual singers in the choir. Through history, conductors, teachers, and

critics have relied on their own interpretation of pitch and harmony, while listening and/or evaluating a choir. In recent years, a few automatic analysis and evaluation systems have been proposed [1, 2] to provide an informed analysis of choirs in terms of intonation. In general, these systems require the extraction of accurate pitch contours for individual vocal tracks, which has hitherto been a roadblock for analysis, as multi- f_0 extraction systems are not able to provide sufficient pitch precision and accuracy to drive analysis systems from full mixed choir recordings. This can primarily be pinned down to the fact that in a choral recording, multiple singers with similar timbres are singing in harmony, and even the same notes within each choir section, leading to overlapping harmonics, which are difficult to isolate. While several multi- f_0 estimation systems have been designed for music with easily distinguishable sources, e.g. music with vocals, guitar, bass and drums, very few research has been carried out in the domain of vocal ensembles, be it because of the lack of annotated datasets or because modeling several people singing very close frequencies, i.e. in unison, is very challenging in terms of f_0 resolution.

In this work we address the computational modeling of pitch in choir recordings. In order to do that, we first evaluate how a set of multi- f_0 estimation algorithms perform with vocal quartets and try to identify their main limitations. Then, we use the evaluation results to select the best-performing algorithm and use it to extract a first approximation of the f_0 of each choir section. In the second step we use a set of traditional DSP techniques to increase the pitch resolution around the estimated f_0 s and model f_0 dispersion. The main focus of this adaptation is not to obtain an accurate f_0 estimate for each voice inside each choir section, but to model the distribution of f_0 of a choir section singing in unison, measured through the dispersion of pitch values across each part.

The rest of the paper is organized as follows: Section 2 provides a brief overview of the current state-of-the-art for multi- f_0 extraction. Section 3 describes the limitations of current systems to characterize f_0 in unison performances. Then, in Section 4 we define the evaluation metrics commonly used in the field, followed by Section 4.1 presenting the dataset used in this study. Section 5 discusses the initial evaluation of state-of-the-art methodologies on our particular material. Following this, Section 6 presents a novel approach to model unison recordings by combining a multi- f_0 estimation algorithm with traditional DSP techniques. Section 7 presents and discusses the results and

limitations of the proposed system, and finally in Section 8 we provide some conclusions on the method and comments on future research that we intend to carry out.

2. STATE OF THE ART

Multi- f_0 estimation involves the detection of multiple concurrent f_0 from an audio recording [3] and it is a core step of the task of automatic music transcription (ATM): converting an acoustic musical signal into some form of musical notation [4]. We briefly summarize a set of multi- f_0 estimation methods that can be applied to vocal music, although they try to estimate f_0 values of individual singers, while we address the modeling of f_0 of unison singing.

Duan et al. [5] presented an approach to multi- f_0 estimation using maximum-likelihood, where they model two different regions of the input power spectrum: the peak region, comprising the set of frequencies that are within a distance d of the peak frequencies, and the non-peak region, which is the complement of the peak region. The input signal is normalized and the power spectrum is computed frame-wise. A set of spectral peaks are extracted using a peak detection algorithm, and then several f_0 candidates are computed in the range of one semitone around each peak. For each time frame, the f_0 of each source is estimated by maximizing the probability of having harmonics that explain the observed peaks and minimizing the probability of having harmonics in the region where no peaks were observed. This is accomplished optimizing the parameters of a likelihood function that combines the peak region likelihood and the non-peak region likelihood, treated as independent sets. They use monophonic and polyphonic training data to learn the model parameters. Their system also estimates polyphony, i.e. how many sources there are in the mix, which define the number of f_0 the model should estimate at each frame. Finally, a post-processing step using information for neighbouring frames is implemented to make the pitch predictions more stable. This process of refining the f_0 estimates, however, removes duplicate estimates, which is a problem in the case of several sources producing the same f_0 , i.e. unison singing. The system parameters were learned using training data consisting of mixes of individual monophonic note recordings from 16 instruments including flute, saxophone, oboe, violin, bass, and violin among others. Then, they evaluated the algorithm on 4-part Bach chorales performed by a quartet: violin, clarinet, tenor saxophone and bassoon. These details about the data they used to train and evaluate the system are very relevant for our research, since given the lack of vocal data in the training and evaluation stages, we expect the system to perform worse in choral music.

Another relevant system for multiple f_0 estimation is the one developed by Klapuri [6], which estimates each f_0 in the mixture iteratively: at every step, the system detects the most predominant f_0 and its corresponding harmonics are then subtracted from the spectrum. In this case, the input signal is passed through a bank of linear band-pass filters that resemble the inner ear behaviour in terms of frequency selectivity. Then, the output signal at each band is processed in a nonlinear manner to approximate

the firing activity in the auditory nerve. After this pre-processing steps, a frequency-domain signal representation is obtained by combining the band spectra, which is then used to extract a pitch salience function to emphasize the fundamental frequencies present in the signal. From this representation, multiple f_0 values are estimated iteratively: at each step, a f_0 value is estimated and its harmonic partials are removed from the spectrum. This step is repeated until a f_0 value is estimated for each source. In [6] he also implements a polyphony estimator, which determines how many f_0 values need to be extracted and therefore the number of iterations. The system was evaluated with a collection of mixtures of individual sounds (some of them from the same source as in [5]). The authors does not explicitly mention any vocals in the dataset, and therefore we assume the method is not optimized for our particular material.

Schramm and Benetos [7] presented a method specifically designed for multi- f_0 estimation in a *cappella* vocal ensembles. They use a two-step approach: first, a system based on probabilistic latent component analysis (PLCA) employs a fixed dictionary of spectral templates to extract the first frequency estimates; as a second step, a binary random forest classifier is used to refine the f_0 estimates based on the overtones properties. Spectral templates are extracted from recordings of multiple singers singing pure vowels in English. These recordings belong to the RWC dataset [8]. This method uses the normalized variable-Q transform (VQT) as input, which is then factorized using the expectation-maximization algorithm to estimate the parameters of the model. As opposed to the previously presented methods, this one is focused on vocal ensembles and it is trained and evaluated with such data, i.e. vocal quartets, one singer per part. They use a f_0 resolution of 20 cents, which is enough for transcription purposes but not to deal with frequencies as close to each other as in unisons.

Another method for multi- f_0 estimation designed for the case of multiple singers is the one presented by Su et al. [9]. The authors claim that data is crucial to develop and evaluate such systems, and yet there is not a labeled multi- f_0 dataset for choir, which is one of the most common type of music through the ages and cultures. Their work has two separate parts: first, they present a novel annotated dataset of choir and symphonic music; then, they build an unsupervised approach to multi- f_0 estimation using advanced time-frequency (TF) analysis techniques such as the concentration of time and frequency method. According to their paper, these techniques help improving the stabilization of pitch, which is interpreted in three dimensions: frequency, periodicity, and harmonicity.

Recent advancements in deep learning based systems have led novel deep learning based multi- f_0 extraction systems, designed to be agnostic to the exact source of the pitched instruments in the mix. *DeepSalience* [10] is one of the most recent systems for multi-instrument pop/rock songs and mixtures. The model leverages harmonic information provided by a HCQT transform, comprising 6 constant-Q transforms (CQT), with a convolutional neural network (CNN) to extract pitch salience from an input audio mixture. The network is fully convolutional with 5 convolu-

tional layers, it uses batch normalization and rectified linear units (ReLU) at each output. The final layer of the network uses logistic activation, mapping each bin of the output to the range [0,1], representing pitch salience. It is trained using cross-entropy minimization. This pitch salience essentially predicts the probabilities of the underlying pitches being present in the input signal with a resolution of 20 cents. Then, using this salience intermediate representation, they use a threshold to estimate multiple frequencies at each frame.

These methods are capable of extracting multiple f_0 from a great variety of audio signals, including music and speech. Most of them are designed for polyphonic signals where each melody is produced by a single source: one instrument or singer. However, the subject of our study are choirs, which involve unison ensemble singing, i.e. performances where several people sing the same notes. Unison recordings are challenging for multi- f_0 estimation because of the possible imprecision in the pitch produced by multiple singers or musicians [9]. Since we focus on the analysis and synthesis of choral singing, it is crucial to take this aspect into account to build models that consider these pitch imprecision.

3. PROBLEM DEFINITION AND APPROACH

The characterization of pitch distribution in unison and choir singing has not been widely studied. Most of the research in this topic is authored by Sundberg [11] and Ternström, who published a review on choir acoustics [12] and carried out several experiments to study f_0 dispersion in unison singing [13]. The authors define f_0 dispersion as the small deviations in f_0 between singers that produce the same notes in a unison performance. This magnitude is directly related to the *degree of unison*, which Sundberg defines as the agreement between all the voices sources. In a later work, Jers and Ternström [14] measured the dispersion between singers and found it to range between 25 and 30 cents.

In multi- f_0 estimation systems, we usually focus on the extraction of a single pitch per source, and state-of-the-art algorithms would then provide a f_0 value for each choir section. However, several singers produce slightly different values in each of the voices of a choir. Then, the question of which is the correct value to be estimated arises: most multi- f_0 estimation algorithms do not have enough resolution to discern the individual pitches, which leads to a potentially imprecise estimation. This suggests that unison performances need to be treated in a different way. Ternström [13] claims that while solo singing has tones with well-defined properties, i.e. pitch, loudness, timbre, unison ensemble singing has tones with statistical distributions of these properties, and we need to consider those when modeling them.

In a recent study, Cuesta and al. [1] created the Choral Singing Dataset (see Section 4.1) to analyze f_0 dispersion in unison ensemble singing by modeling the distribution of fundamental frequencies. Using individual tracks for each singer, they extracted f_0 curves and computed the mean f_0 as the perceived pitch and the standard deviation of the

distribution as the f_0 dispersion. In Figure 1 we display an example of the f_0 trajectories of four sopranos, where we observe that there are slight f_0 differences between them. This study found dispersion values ranging from 20 to 30 cents on average, depending on the choir section and the song, which agrees to previous literature [13]. However, this type of analysis requires an individual audio track for each singer in the mixture, and this data is difficult to obtain given that choirs are not recorded using this set up. This particular limitation leads us to explore in the present study ways of analyzing choir recordings directly from the singer mixture, which involves dealing with four different melodies (SATB), each of them involving a unison.

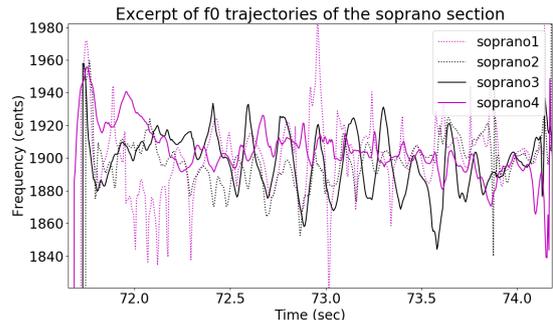


Figure 1: F_0 curves of four sopranos singing the same note. We see how the curves oscillate and differ from each other.

In this study, we propose a methodology for pitch content analysis on unison ensemble singing that has two main stages:

1. **Multi- f_0 estimation.** In the first stage, we perform multi- f_0 estimation in the audio mixture in order to roughly estimate the pitches of the four voices of the choir. For this part, we evaluate the performance of a set of existing multi- f_0 estimation systems in the context of vocal quartets, where we have precise f_0 ground truth information, to select the one with a better performance.
2. **F_0 -dispersion modeling** In the second stage, we refine the frequency analysis around those pitches to further characterize f_0 dispersion in each of the unison voices. In order to do so, we consider a DSP-based approach and adapt a method with higher frequency resolution to model each melodic source as a distribution of f_0 s instead of a single value.

4. EVALUATION METHODOLOGY

As mentioned in previous sections, we evaluate the performance of three state-of-the-art algorithms for multi- f_0 estimation in vocal quartets, e.g. SATB with a single singer per section, in order to investigate which method is more suitable for this music material. We consider the methods proposed by Klapuri (KL) [6], Schramm et al. (SCH) [7] and Bittner et al. (DS) [10], all of them publicly available and representative of the state of the art in the area.

4.1 Dataset

There are very few datasets of choral music which are annotated in terms of f_0 . In our experiments, we take advantage of the Choral Singing Dataset¹ further described in [1]. This dataset was recorded in a professional studio and contains individual tracks for each of the 16 singers of a SATB choir, i.e. 4 singers per choir section. Although each section was recorded separately, synchronization between all audio tracks was achieved using a piano reference and a video of the conductor that singers followed during the recording.

This dataset comprises three different choral pieces: *Locus Iste*, written by Anton Bruckner, *Niño Dios d'Amor Herido*, written by Francisco Guerrero, and *El Rossinyol*, a Catalan popular song; all of them were written for 4-part mixed choir. This dataset is more suitable for this study than the one presented in [9]: having the individual tracks of each singer allows us to create *artificial* mixes between voices, e.g. vocal duets or quartets, small choir, large choir...etc. Using different combinations of all 16 singers, we created 256 SATB quartets for each piece, which represent all possible combinations of singers taking into account the voice type restriction, i.e. we need one singer per voice. These vocal quartets are used to evaluate the performance of the three algorithms.

4.2 Multi- f_0 evaluation metrics

In multi- f_0 estimation systems, there are multiple f_0 values per frame n . Following the terminology used by Bittner [15], we define the ground truth value(s) in frame n as $f[n]$ and the estimation as $\hat{f}[n]$, which denote the pitches of all active sources in that frame.

For a given frame n we denote as true positives, $TP[n]$, the number of correctly transcribed pitches, and as false positives, $FP[n]$, the number of pitches present in the estimation, $\hat{f}[n]$, which are not present in the ground truth, $f[n]$. Similarly, the false negatives value, $FN[n]$, measures the number of pitches present in the ground truth which are not present in the estimation. Based on these, we define the following set of metrics: *accuracy*, *precision* and *recall*, and a set of errors: the substitution error (E_{sub}), miss error (E_{miss}), and false alarm error (E_{fa}). Finally, total error, E_{tot} is reported as the combination of E_{sub} , E_{miss} and E_{fa} .

All the presented evaluation metrics also have their associated *chroma* versions, which considers an estimated f_0 to be correct if it is one octave apart from the corresponding target pitch. For more details about these metrics we refer the reader to [15].

5. MULTI- F_0 ESTIMATION RESULTS

All the SATB quartets of the dataset were evaluated in terms of multi- f_0 estimation: by means of the individual tracks, we extracted f_0 curves for every singer using the spectral-amplitude autocorrelation (SAC) method [16] and

we then combined them to create the multiple f_0 ground truth at each frame.

A summary of the results is displayed in Figure 2, where we present the accuracy, recall and precision averaged for each of the algorithms in the three songs of the dataset. We observe that *DeepSalience* (DS) outperforms Klapuri (KL) and Schramm (SCH). It is also interesting to point out that the difference between these metrics and their *chroma* versions is very small, thus suggesting that the three algorithms are fairly robust in terms of octave errors. We also observe that the algorithm by Schramm et al. has a higher variability with respect to the other ones, suggesting that its performance is highly dependent on the input signal. Also, it is important to mention that while KL and DS predict multi- f_0 values from a long audio file, i.e. a full choral recording, we splitted our audio material in shorter clips (each of them 10 seconds long) to evaluate SCH method: the PLCA algorithm employed in this method is computationally very expensive and we could not obtain results using the full recordings.

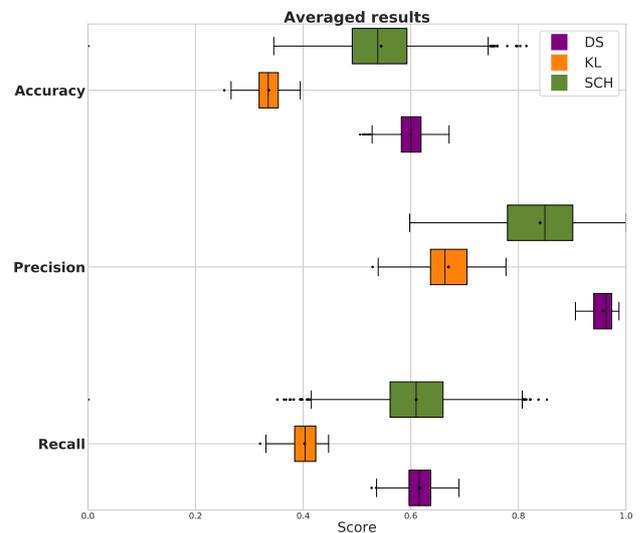


Figure 2: Accuracy, precision and recall for each of the three algorithms (DS, DU, SCH) averaged over all the dataset, i.e. all the SATB quartets.

In terms of error analysis, Table 1 provides the average errors for each of the algorithms. These results suggest that extracting multiple frequencies from a vocal ensemble is a very challenging task, since the total error is almost 40% in the best performing method. A part from this, we can extract a few more insights: all algorithms have a very low false alarm error, which means that they almost never report an f_0 when there is not one in the ground truth; in addition, *DeepSalience* does a good job regarding the substitution error, which means that it rarely reports a wrong f_0 . However, the miss error is pretty high, especially in Klapuri's algorithm, which means that there are a lot of f_0 that are not extracted. In the case of Schramm's method, though, the miss error is lower, suggesting that their voice assignment step improves the performance of the multi- f_0

¹Choral Singing Dataset: <https://zenodo.org/record/1319597>

	DS	SCH	KL
Substitution error	2.3%	10%	12%
Miss error	35%	28%	48%
False alarm error	0.4%	1.5%	8%
Total error	38%	40%	67%

Table 1: Summary of error metrics in multi- f_0 estimation .

estimation. We could also relate these differences to the fact that SCH is the only method designed for and trained with singing voice data. However, given the length limitation of SCH and based on the overall results, we select *DeepSaliency* as the method to be used in the first stage of our model.

6. F_0 DISPERSION MODELING

The first step of our method presented above uses *DeepSaliency* to extract multiple pitch estimations at each frame of the audio input. In the ideal case, at this stage we would obtain one f_0 value for each choir section; however, as discussed in Section 5, although *DeepSaliency* is the best-performing algorithm from the evaluated set, there are still some errors in the output.

In the second stage of our method we consider traditional DSP techniques to increase the frequency resolution of our model. We compute the spectrogram of the input audio signal using a Hanning window of 4096 points zero-padded to twice its length, resulting in an FFT size of 8192. An excerpt of this spectrogram is displayed in Figure 4 with magnitude in dB and the frequency axis in cents, where we observe that f_0 values for each choir section are well-separated.

With this time-frequency representation, we then locate each of the estimated fundamental frequencies (*DeepSaliency* output), which will ideally match one of the spectral peaks. Even though we use a large FFT size, since we want to obtain a high pitch resolution, we interpolate the peaks and recompute the peak locations as the maximum value of each interpolated peak. This process is illustrated in Figure 3, where the top and bottom plots correspond to a vocal quartet and full choir spectrum, respectively. The dashed black line represents the original spectrum, while the red solid lines and the green asterisk correspond to the interpolated peaks. We observe that the peaks in the full choir case (bottom) have less energy and are a bit more noisy than the vocal quartet ones (see third and sixth peaks for example).

Once we have this information, we compute the bandwidth of each peak as a measure of the dispersion of the f_0 distribution in the unison case. Remember that our aim is to characterize the distribution of f_0 for each choir section rather than obtaining a single f_0 value. For each choir section, we find and interpolate the peak in the spectrum and consider the peak frequency as the mean frequency of the distribution and its bandwidth as its dispersion. The bandwidth is expressed in cents (computed with a refer-

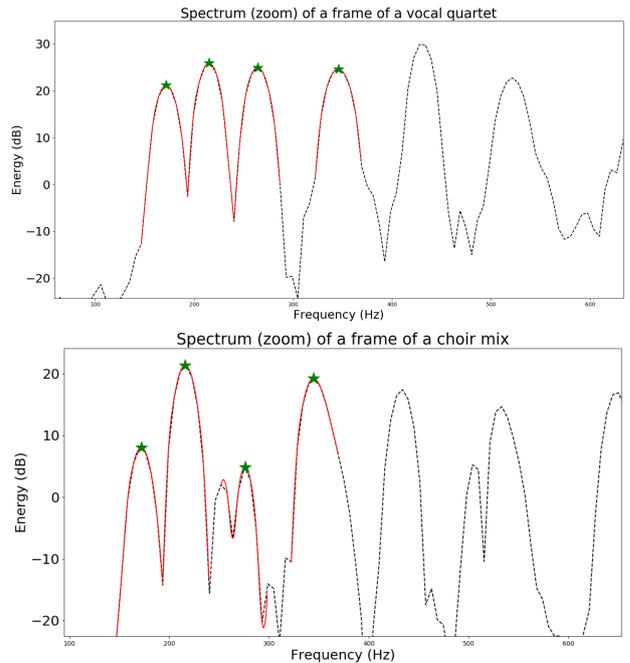


Figure 3: Example frames of the spectrum with the interpolated peaks corresponding to the estimated f_0 s (green asterisk). The top plot corresponds to a 4-singers audio input (quartet), while the bottom plot corresponds to the 16-singer audio input (full choir).

ence frequency of 220 Hz) and computed as follows:

$$f_{0_{dispersion}} = b_2 - b_1 \quad (1)$$

where b_2 and b_1 are the frequency bins around the spectral peak where the amplitude of the spectrum decays 3 dB. Note that in this first approach we do not take into account the window type and size used in the spectral analysis, although they influence the peak bandwidth. In further studies, we plan to study and document the effect of these two analysis parameters in the dispersion computation.

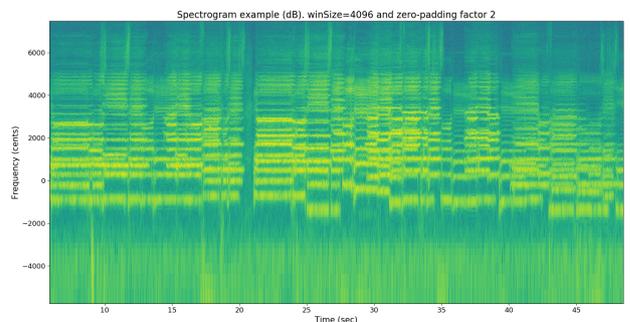


Figure 4: Spectrogram (zoomed) of the piece Locus Iste computed using a Hanning window and $N = 4096$ zero-padded to twice its length, resulting in an FFT size of 8192.

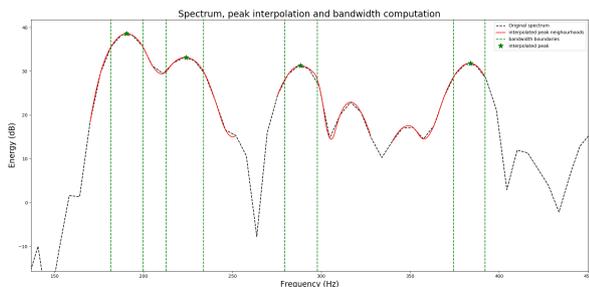


Figure 5: Example frame of the spectrum with the interpolated peaks and the boundaries for the bandwidth computation.

7. RESULTS

In this Section we present the results of the dispersion analysis averaged by piece and by choir section. Since the first step of the presented framework outputs noisy multi- f_0 estimations, the results of the second step displayed here are obtained using the ground truth pitch values to locate the peaks. This allows us to perform a better evaluation of the dispersion computation method, since the errors that come from the multi- f_0 estimation are not present here. In the real application, the ground truth pitches would be replaced by the output of the selected multi- f_0 estimation algorithm. Figure 5 shows another example frame of the analysis, with the peak interpolation and also vertical lines showing the bandwidth delimitation.

Table 2 provides the f_0 dispersion results averaged by choir section and piece. *BTAS* stand for *bass, tenor, alto, soprano*, and *Q* and *CM* are short for *quartet* and *choir mix*, meaning that the dispersion values belong to a 4-singers and 16-singers setting, respectively. We would expect the dispersion values to be larger in the 16 singers case, because having several voices producing similar frequencies might generate wider spectral peaks. The differences on average are not very strong; however, we conducted an independent-samples t-test to compare the performances of vocal quartets with the full choir and found that the differences were significant. For example, there was a significant difference in the bass section of the quartet singing *El Rossinyol* ($M=181, SD=39$) and the bass section in the choir mix ($M=191, SD=79$), $\alpha = 0.05, p < 0.001$. This tendency applies to the whole dataset, although the effect size is small (around 0.2 on average) and therefore we might need a deeper analysis to find out if the differences are not only statistically significant but also relevant in terms of pitch content.

These results can not be directly compared to any ground truth, since previous studies modeled pitch dispersion in different ways, i.e. standard deviation of the distribution [1] or bandwidth of the partial tones [13]. Instead, we compare the tendency of the results with the ones obtained by the authors in a previous study where individual pitch tracks, and thus accurate individual f_0 estimations, were used [1]. In Table 3 we display the results from the men-

	Locus Iste		El Rossinyol		Niño Dios	
	Q	CM	Q	CM	Q	CH
B	231 (57)	248 (130)	181 (39)	191 (70)	227 58	257 175
T	136 (30)	140 (38)	132 (26)	136 (30)	143 (31)	149 (40)
A	100 (22)	104 (25)	98 (19)	103 (23)	105 (22)	110 (28)
S	76 (23)	79 (27)	75 (16)	80 (20)	78 (20)	82 (25)

Table 2: Average f_0 dispersion results computed as the bandwidth of the peaks in the whitened spectrum. Dispersion values are in **cents**. *Q* refers to a SATB quartet with one singer per section and *CM* refers to a SATB choir mix with 4 singers per section. Values in italics are standard deviations also in cents.

Soprano	Alto	Tenor	Bass
20.16	22.66	22.22	26.02
Locus Iste		El Rossinyol	Niño Dios
19.32		27.65	21.33

Table 3: Averaged results of pitch dispersion from [1]. All values represent dispersion in cents, computed as the standard deviation of the pitch distribution.

tioned work, where the dispersion is averaged by choir section (all three pieces together) and by piece (all four choir sections together). These results suggest that the dispersion (in cents) is higher in the bass section of a choir, which is also true for our results. Following basses, the presented results have tenors, followed by altos and finally sopranos, with the lowest average dispersion. In Table 3 we have altos and tenors with very similar average values, and sopranos also obtained the lowest dispersion values. Therefore, although the dispersion magnitude can not be compared, the trend is very similar, suggesting that the analysis is consistent.

After an analysis of the results, we observe that this framework has a few limitations, including that the results highly depend on the performance of the multi- f_0 estimation algorithm employed in the first stage. In this paper we evaluated three algorithms which are considered state-of-the-art and *DeepSalience* was selected according to the evaluation we carried out, but we hypothesize that using a system specially designed for singing voice, and even for choral singing, e.g. [7] and [9], might improve the performance of the whole method: if the f_0 estimates at each frame are precise, then the peaks would be modelled correctly (as happened in the dispersion evaluation), yielding more accurate f_0 dispersion values. However, these methods were dismissed from the final approach because of the length limitation [7] and because it is not publicly available [9].

The proposed framework models the f_0 distribution of a unison using two values: the f_0 , extracted in the first stage and refined in the second stage, and the f_0 dispersion. We

believe these values are a good descriptor for unison performances, but a more complex model incorporating temporal evolution analysis could also be considered and used to estimate the quality of a choir or in choral synthesis applications.

8. CONCLUSIONS

In this paper we presented a framework for the f_0 modeling in choral singing recordings that uses a two-stage methodology: first, a deep learning based multi- f_0 estimation system is employed to obtain one pitch value for each choir section; second, we locate these frequencies in a whitened version of the spectrum of the input signal with a higher pitch resolution. This process allows us to model each unison as a distribution of f_0 , characterized by two values: the mean f_0 and the f_0 dispersion. The preliminary results we obtained do not show strong differences in the dispersion between a large and a small number of singers, but more data might reveal different patterns. Since we evaluated this framework as a case study, more experiments will be done in the near future, and a deeper analysis of the relationship between the time-frequency representation and the results will be carried out to make the system more effective.

Acknowledgments

This work is partially supported by the European Commission under the TROMPA project (H2020 770376) and by Spanish Ministry of Economy and Competitiveness under the CASAS project (TIN2015-70816-R). First author is supported by FI Predoctoral Grant from AGAUR (Generalitat de Catalunya). We thank Rodrigo Schramm and Andrew McLeod for the support running their code with our audio material.

9. REFERENCES

- [1] H. Cuesta, E. Gómez, A. Martorell, and F. Loáiciga, "Analysis of Intonation in Unison Choir Singing," in *Proc. of the 15th International Conference on Music Perception and Cognition (ICMPC)*, Graz (Austria), 2018.
- [2] J. Dai and S. Dixon, "Analysis of Interactive Intonation in Unaccompanied Satb Ensembles," in *18th International Society for Music Information Retrieval Conference*, 2017.
- [3] A. McLeod, R. Schramm, M. Steedman, and E. Benetos, "Automatic Transcription of Polyphonic Vocal Music," *Applied Sciences*, vol. 7, no. 12, p. 1285, 2017. [Online]. Available: <http://www.mdpi.com/2076-3417/7/12/1285>
- [4] 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.
- [5] 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.
- [6] A. Klapuri, "Multiple Fundamental Frequency Estimation by Summing Harmonic Amplitudes," in *Proc. of the International Society for Music Information Retrieval Conference*, Victoria, Canada, 2006, pp. 216–221.
- [7] R. Schramm, E. Benetos *et al.*, "Automatic Transcription of A Cappella Recordings from Multiple Singers." Audio Engineering Society, 2017.
- [8] M. Goto, H. Hashiguchi, T. Nishimura, and R. Oka, "RWC Music Database: Music Genre Database and Musical Instrument Sound Database," 2004.
- [9] L. Su, T.-y. Chuang, and Y.-h. Yang, "Exploiting Frequency, Periodicity and Harmonicity Using Advanced Time-Frequency Concentration Techniques for Multi-pitch Estimation of Choir and Symphony," in *Proc. International Society for Music Information Retrieval Conference*, 2016.
- [10] R. M. Bittner, B. McFee, J. Salamon, P. Li, and J. P. Bello, "Deep Saliency Representations for F0 Estimation in Polyphonic Music." in *Proc. of the International Society for Music Information Retrieval Conference*, New York, 2017, pp. 63–70.
- [11] J. Sundberg, *The Science of the Singing Voice*. Northern Illinois University Press, 1987.
- [12] 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.
- [13] ———, "Perceptual evaluations of voice scatter in unison choir sounds," *STL-Quarterly Progress and Status Report*, vol. 32, pp. 041–049, 1991.
- [14] H. Jers and S. Ternström, "Intonation analysis of a multi-channel choir recording," *TMHQPSR Speech, Music and Hearing: Quarterly Progress and Status Report*, vol. 47, no. 1, pp. 1–6, 2005.
- [15] R. Bittner, "Data-driven fundamental frequency estimation," Ph.D. dissertation, New York University, 2018.
- [16] E. Gómez and J. Bonada, "Towards Computer-Assisted Flamenco Transcription: An Experimental Comparison of Automatic Transcription Algorithms as Applied to A Cappella Singing," *Computer Music Journal*, vol. 37, no. 2, pp. 73–90, 2013.