

# Optimizing Melodic Extraction Algorithm for Jazz Guitar Recordings Using Genetic Algorithms

**Sergio Giraldo**

Pompeu Fabra University  
sergio.giraldo@upf.edu

**Rafael Ramirez**

Pompeu Fabra University  
rafael.ramirez@upf.edu

## ABSTRACT

Extraction of the main melody of a musical piece is a preliminary step in the process of transcribing the piece. Automatic melodic extraction is the task of computationally extracting what a human listener would perceive as the main melody of a polyphonic recording. Several melodic extraction systems have been proposed. However, such systems normally require a number of parameters to be manually tuned in order to accurately perform melody extraction in different contexts, i.e. instruments combinations. In this study, we propose a methodology for automatically optimizing some relevant parameters of a melody extraction algorithm using genetic algorithms. We simultaneously obtained both MIDI and audio recordings of jazz standards, and we collected commercial audio recordings extracted from jazz guitar CDs. Based on the MIDI recordings as ground truth, two different instrument settings are compared (Jazz trio and quartet), as well as different audio mixing of the melody with respect to the accompaniment track. We show that, compared to using the default parameters, the overall accuracy of the melody extraction with the optimized parameters is improved.

## 1. INTRODUCTION

In jazz education, the extraction of the main melody of musical piece is a "must to do" task. In contrast to classical music, in popular music melody extraction is the way to understand melodic variations or performance actions that lead to play in certain style. Similar to spoken language, popular music is often taught mainly by oral tradition, thus the student learn by listening and copying experts' performances. The task of automatically extracting the main melody of a polyphonic recording is a computational problem widely studied in recent years. Given a recording of a polyphonic musical piece (performed either by several instruments or by a polyphonic instrument), the task is to computationally identify the main melody of a piece. The main melody may be defined as the monophonic sequence of pitches that a listener would hum (or whistle) when listening to a polyphonic piece of music regardless the music style. Several applications of automatic

melody extraction can be mentioned, for instance query by humming [1], soloist identification [2], and automatic accompaniment [3].

In this study we focus on the salience pitch approach method used by Salamon et al. *Melodia* [7], in which melodic extraction is performed in several steps: sinusoidal extraction, salience function computation, pitch contour computation and characterization, and melody selection. For the computation of each of these steps different parameters such as minimum and maximum frequency, peak distribution and peak frame threshold, pitch and time continuity need to be defined. In *melodia* system, these parameters have been tuned using the grid search method over different data bases in which melody was performed mainly by singing voice. However, the default parameters' values fail to extract the melody for some specific instrument settings (jazz guitar recordings in our case). We propose a method to optimize the algorithm's main parameters using genetic algorithms, based on specific instrument settings. In this study we focus on two instrument settings: Trio setting consisting of electric guitar, bass and drums, and Quartet setting consisting of electric guitar, bass, drums and piano. In both cases the melody is performed by the guitar.

The rest of the paper is organized as follows: Section 2 presents the background to automatic melodic extraction, a brief description of the *Melodia* algorithm, and a brief introduction to genetic algorithms. Section 3 describe our data set of songs and presents our methodology for optimizing the algorithm parameters. In Section 4 we present our results, and finally, Section 5 presents our conclusions and future work.

## 2. BACKGROUND

### 2.1 Melody extraction

The most widely used methods for melodic extraction are the ones based on salience pitch calculation. Usually these methods are performed in three steps. First, a spectral representation of the signal is computed using spectral analysis techniques (e.g. Fast Fourier Transform). Second, a salience function (time-frequency representation of pitch salience) is computed to obtain several  $f_0$  candidates, usually using weighted harmonic summation methods. Finally, the melody peaks are selected based on tracking methods over frequency and time. Other methodologies use source separation based on timber models and grouping principles [4, 5], and some use stereo separation to estimate the panning of each source [6].

In this study we focus on *Melodia* [7]. In their approach melodic extraction is performed by several steps. First, *sinusoidal extraction* filters the signal to enhance the most audible frequencies of human hearing range. Then, the Short Fourier Transform is calculated using a Hanning window of 46.4ms with a hop size of 2.9ms and a 4 zero padding factor. Finally, the peaks are corrected using instantaneous frequency method [8]. In a second step, the *saliency function computation* is performed based on the summation of the weighted energy of the harmonic peaks of a given frequency in order to obtain the  $f_0$  candidates. The number of harmonics considered and the weighting scheme is an important factor that affects the saliency computation. The third step is *pitch contour computation*, in which the peaks detected in the previous step are grouped into pitch contours, based on several thresholds defined per frame basis, as well as per time continuity basis. Some of the algorithm's key parameters in this step are the following:

- Peak Distribution Threshold: Allowed deviation below the peak saliency mean over all frames (fraction of the standard deviation)
- Peak Frame Threshold: Per-frame saliency threshold factor (fraction of the highest peak saliency in a frame)
- Pitch Continuity: Pitch continuity cue (maximum allowed pitch change during 1 ms time period)
- Time Continuity: Time continuity cue (the maximum allowed gap duration for a pitch contour)

The fourth step is *contour characterization* in which the melody contour is chosen among the contours created in the previous step. To do so, a set of features that guide the system to select the main melody is implemented. These features include pitch deviation, pitch trajectory, presence of vibrato, as well as contour pitch, length and saliency. Finally the last step is *melody selection* which is performed by voicing detection, octave error minimization, pitch outliers removal, and final melody selection. A complete explanation of each step can be found in [7].

## 2.2 Genetic Algorithms

Genetic Algorithms (GA) are stochastic optimization algorithms which imitate the biological mechanisms of natural selection. GA are widely used in several optimization problems involving discontinuous, noisy, high-dimensional, and multi-modal objective functions. In contrast to other optimization algorithms, the search that GA's perform is done in a more global context, whereas others (e.g. gradient descent) perform search in a more local context. Genetic algorithms work as follows: First, an initial random population of individuals is created. This initial population serves as seed to generate future generations by means of combining different parents (crossover) and randomly modifying a single individual (mutation). The algorithm iterates, and in each iteration it selects the best subjects

based on a fitness function. The next generation of individuals is generated by applying mutation and crossover. The individuals with best fitness values are preserved in the next generation, while the others are discarded. The process iterates until a maximum number of iterations is reached or the fitness relative value of the best individual does not change more than a tolerance value during several generations. The most popular applications of GA in a musical context have been done for music composition [9, 10]. Other approaches have used GA's for automatic music transcription [11], music segmentation [12], and interactive music applications [13], among many others.

## 3. METHODOLOGY

Two experiments were proposed to test our methodology for optimizing melody extraction. A first experiment was done using 22 simultaneous audio midi recordings. Ground truth was generated from the recorded MIDI data, and mixes were created using the synthesized version of the MIDI recording. The second experiment was done in a more realistic context using four recordings in which melody is recorded separately from the accompaniment track. In this case ground truth was build from automatic extraction of the pitch contour of the melody track, using YIN algorithm [14] and performing manual correction afterwards.

### 3.1 First Experiment Data

The train data set consisted of 22 jazz standard tunes, recorded by a professional guitarist. The indications given to the musician were that melodies should be played monophonic, with no chord strums, or double notes. Only the main melody of the tune was recorded. Non repeated sections of the recorded melody were selected. For example, if a tune had the form AABA, only the second A and the B parts were used. The melodies were recorded over commercial backing tracks recordings, in which bass and piano parts are recorded separately on each stereo channel [15]. This made possible to mix different instrument settings: drums, bass and guitar (trio) and piano, drums, bass and guitar (quartet).

#### 3.1.1 Building ground truth

All the melodies were simultaneously recorded in both audio and MIDI, using a commercial guitar to MIDI converter (Sonnus). Because the accuracy of the device is not perfect, manual correction of the audio to MIDI converted melody was performed after each recording. Each MIDI event (note) has information of pitch (midi number), onset (in seconds) and duration (in seconds). Using this information and a hop size of 46.4ms (as used by the melody extractor algorithm) it was possible to obtain the pitch contour of the melody frame by frame. Finally, to assure that the pitch contour corresponds exactly to the melody to be extracted from the created mix, the MIDI sequence of the melody was synthesized in wav format (using the built-in guitar synthesizer of the DAW - Digital Audio Workstation) and mixed with the accompaniment backing track.

### 3.1.2 Experiment overview

Our aim was to find a set of parameters that work well for tunes of similar type. We defined two main groups of tunes: one was an audio mix of the recorded tunes with synthesized guitar, bass and drums (trio set), and the other was the audio mix of the same recorded tunes with synthesized guitar, bass, piano, and drums (quartet set). For each of the above sets we followed a similar method as [7], creating different audio mixes in which the melody is at three different sound levels (-6dB, 0dB and +6dB) with respect to the overall sound level of the backing track. Eighteen of the tracks were used as train data to optimize parameters. The optimized parameters were tested in the four songs that were left aside for testing.

## 3.2 Second Experiment Data

In a second experiment, the data set consisted of four standard jazz tunes obtained from a guitar teaching series book [16] in which melody is recorded separately from the accompaniment track. This set was chosen to test our optimization strategy in a more real context in which common effects such as reverb can make melody extraction more difficult.

### 3.2.1 Building ground truth

In this case, the pitch contour of the melody was obtained from the melody track using Yin algorithm [14]. We set the frame resolution for extraction of the algorithm to 46.4ms to be consistent with the frame resolution of *Melodia*. However, manual correction needed to be performed. To do so, we implemented a simple GUI (guided user interface) to plot the obtained pitch contour over the spectrum of the melody. This facilitated the visualization and correction of errors such as octave shift and/or missing pitches.

### 3.2.2 Experiment overview

We set up different mixes with different sound levels of the melody, similar to experiment number one. For each mix we used the method of "leave one out", so we used three of the song to optimize parameters, and the fourth one to test the optimized parameters. This process was repeated four times leaving each time a different song, so all the songs were used for training, as well as for testing.

## 3.3 GA's for Melody Extraction Optimization

As explained in Section 2 *Melodia* algorithm uses several steps in the extraction of the pitch contour. Each of these steps uses different parameters and thresholds for the calculation. Each of these parameters were reported to have been calculated using grid search, over different test sets [7], in which the main melody was mainly vocals. Our aim in this study is to optimize these parameters for the extraction of the melody played by a clean electric guitar in a jazz guitar context.

### 3.3.1 Parameters

After some initial tests the parameters found to be more sensitive for melodic extraction accuracy were:

- Maximum Frequency
- Minimum Frequency
- Peak Distribution Threshold
- Peak Frame Threshold
- Pitch Continuity
- Time Continuity

### 3.3.2 Fitness function

In order to manipulate the *Melodia* parameters we used the algorithm implementation found in *Essentia Library* [17]. We manipulate the code to set the selected parameters as input variables. A fitness function was implemented in Matlab in which the algorithm is called with some initial parameters (the default ones) and for each parameter a threshold is defined according to the values reported in [7]. The cost of the fitness function is calculated based on the overall accuracy measure reported in [7], which is defined as the total proportion of frames correctly estimated by the algorithm. This proportion is measured based on the true negatives (TN) and true positives (TP) for which the pitch estimation is correct (between a range of +/- 1/4 of tone of the ground truth). TN and TP were calculated as follows:

$$TP = 1200 * \log_2\left(\frac{ExtractedMelody}{GroundTruth}\right) > 50cents$$

$$TP + TN = 1200 * \log_2\left(\frac{ExtractedMelody}{GroundTruth}\right)$$

Thus, the cost (j) was calculated as:


$$j = 1 - \left(\frac{TP}{TP+TN}\right)$$

For the cost calculation all the frames of all the set of melodies were summed.

As an example we present in Figure 1 the extracted contour by *Melodia* algorithm and the ground truth pitch contour extracted from the MIDI data. On the bottom graph differences between both contours are plotted. The main idea behind the optimization process is to minimize the overall errors shown in the graph using the overall accuracy measure explained above.

### 3.3.3 Implementation

In order to optimize the parameters of the melody extraction algorithm we implemented a genetic algorithm using the Matlab implementation of GA [18]. After defining a list of input parameters, a threshold for each parameter and a cost function, the algorithm randomly generates an initial population, in our case it consisted of a set of 20 vectors with different parameter values combinations. The stopping criteria used in this study was set to a maximum of 500 iterations, or a relative threshold change in fitness function of  $1 \times 10^{-6}$ . Crossover factor was set to 0.8 and mutation factor was set to 0.02, which are typical settings for the parameters of genetic algorithms. Upper and lower bounds, as well as initial parameters' values were set as follows:



**Figure 1.** Pitch contour extracted by *Melodia* in a music excerpt (top), and the ground truth pitch contour extracted from the MIDI file (middle). On the bottom we present the detection errors calculated as the difference between the two contours

$$\begin{aligned} \text{Upper Bound} &= [1760, 87, 1, 1, 30, 150] \\ \text{Lower Bound} &= [1000, 55, 0.5, 0.5, 10, 50] \\ \text{Initial Point} &= [1760, 55, 0.9, 0.9, 27.5625, 100] \end{aligned}$$

The values of the vectors correspond to Maximum Frequency, Minimum Frequency, Peak Distribution Threshold, Peak Frame Threshold, Pitch Continuity and Time Continuity, respectively.

#### 4. RESULTS

For the first experiment (18 songs) the results for the training set is presented in Table 1. For each instrument setting (Trio and Quartet) we apply the optimization methodology on the three different mixing levels. The Overall Accuracy was calculated summing all the frames of all the analyzed tunes. The obtained algorithm’s accuracy is presented in the first column of each instrument setting (OADP). After the optimization process the obtained accuracy is presented in the second column (OAAO) for each instrument setting.

	Trio		Quartet	
	OADP %	OAAO %	OADP %	OAAO %
Mix at 0dB	54.25	73.16	46.43	73.49
Mix at -6dB	41.93	55.44	46.34	54.62
Mix at +6dB	79.75	82.80	81.15	82.54

**Table 1.** Overall Accuracy for the training set in experiment one. First column presents the Overall Accuracy with Default Parameters (OADP) and second one the Overall Accuracy After Optimization (OAAO), at different mixing levels for each of instrument settings (Trio and Quartet).

In Table 2 we present the Overall Accuracy obtained in the test set. The melodies of the test songs were extracted using the default parameters, and the optimized ones. The results are shown for both of the instrument settings at the

three different mixing levels. First column presents the Overall Accuracy with Default Parameters (OADP) and second one the Overall Accuracy with Optimized Parameters (OAOP)

	Trio		Quartet	
	OADP %	OAOP %	OADP %	OAOP %
Mix at 0dB	52.29	72.58	59.92	73.16
Mix at -6dB	43.1	53.06	46.95	55.45
Mix at +6dB	81.52	82.80	74.08	81.37

**Table 2.** Overall Accuracy for the test set in experiment one. First column presents the Overall Accuracy with Default Parameters (OADP) and second one the Overall Accuracy with Optimized Parameters (OAOP), at different mixing levels for each of instrument settings (Trio and Quartet).


For the second experiment, we used a set of four songs extracted from a jazz guitar education series book [16]. We optimized the parameters using three of the songs and test them on the extraction of the fourth song (leave one out approach). In table 3 we present, for each fold, the accuracy using the default parameters and the accuracy using the optimized parameters.

	Fold 1		Fold 2	
	OADP %	OAOP %	OADP %	OAOP %
Mix at 0dB	48.41	58.00	49.90	55.86
	Fold 3		Fold 4	
	OADP %	OAOP %	OADP %	OAOP %
Mix at 0dB	56.52	59.57	62.91	68.85

**Table 3.** Overall Accuracy for Test set in experiment two. First column presents the Overall Accuracy with Default Parameters (OADP) and second one the Overall Accuracy with Optimized Parameters (OAOP), at different mixing levels for each leave one out fold.

In figure 2 the error against the number of iterations is plotted for the quartet setting and an mix of +6dBs. In this figure it is possible to notice how after approximately 250 iterations the algorithm finds a minimum. Also, during the first iterations of the algorithm, it is possible to notice how sensitive is the algorithm to the chosen parameters. In table 1 (training data) on each scenario the performance of the algorithm increases after optimization. Also the performance of the algorithm increases as the sound level of the mix increases. The improvement on -6dB scenario is less than in +6dB scenario. This is expected since, as it is the case for human listeners, the melody is harder to be differentiated from the backing track if the sound level is too low. In table 2, the results are similar as the ones described above, but the accuracy is lower in general. This is due to the fact that we are now testing on unseen data. However in all the cases there is a considerable improvement of the

performance using optimized parameters compared to the default ones.



**Figure 2.** Error against number of iterations on scenario quartet and +6db. After 250 iterations a minimum is found

On the second experiment in table 3, it is possible to see how the performance of the algorithm increases in all cases after optimization. Even though the improvement in performance is less than for the MIDI generated melodies, in this scenario we are dealing with more realistic recordings, in which effects, such as reverb, and other artifacts are present. Also hand annotated ground truth may contain human biases, as for example the duration of the notes, or level of legato, may differ between different listeners.

## 5. DISCUSSION

The work presented in this paper refers to a methodology to solve an specific problem, rather than a system to be used to general problems (e.g as a melodic extraction plugin). This same methodology could be applied to other performance settings for other instruments and other music styles. However, the need of an annotated data set will be necessary in any case. For example if we want to improve melody extraction for a certain performance settings (e.g. melody played by saxofone), an annotated data set with the ground truth would be necessary to optimize the parameters of the algorithm for that specific type of musical setting.

Choosing genetic algorithms against other search methods (e.g. Grid Search) was done mainly for practical reasons, as they provide a more efficient strategy when searching over a large dimensional space. If we think of a method such grid search on a six dimensional space (as it is the case in our problem) and we set 10 values per parameter, we would need to try  $10^6$  combinations to extract the melody. Given that in our case each combination the execution of the *Melodia* algorithm, this would much more computationally expensive, compared to the method proposed. In Figure 2 the best fitness value is plotted against the number of iterations for one of the scenarios. A minimum is found after 250 iterations, with a population of 20 individuals per iteration, resulting in 5000 computations. In all scenarios the minimum was found in a similar number of iterations.

## 6. CONCLUSIONS

In this paper we have proposed a method to optimize the extraction parameters of *Melodia*, an algorithm for melody

extraction, using Genetic Algorithms. Our approach is oriented to improve the melodic extraction for specific instrument settings, for problems in which the extraction has to be performed on the same instrument set (eg. electric guitar, bass, piano and drums). We have performed two experiments, a first one using 22 simultaneous midi and audio recordings, in which the MIDI information has been used to build the ground truth. Optimization was performed in 18 songs and 4 songs were left for testing. Two instrument settings were created: trio (guitar, bass, drums) and quartet (guitar, bass, piano, drums). Also different audio mixes were created in which the melody is at different sound levels with respect to the accompaniment track. The second experiment was performed on four commercial recordings in which the melody was recorded in a separated channel from the accompaniment track. In this experiment we used a "leave one out approach" for optimization and testing. Results show that our optimization methodology improves in all cases the overall accuracy of the detection for the first experiment. Improvement in overall accuracy in second experiment is less as we are applying the methodology on a more real context in which aspects like reverberation in the melody track may induce more errors in the melodic detection.

## 7. FUTURE WORK

Although Genetic Algorithms is an approach that provides a computationally efficient strategy to optimize the parameters of the problem, the execution time of our implementation could still be improved. Given our promising initial results, we plan to port our implementation in a faster programming language (eg. C++, or Python). Also it would be interesting to investigate how correlation and mutation factors may affect the results, as well as the number of iterations to minimize the cost function on each of the scenarios.

## 8. REFERENCES

- [1] R. B. Dannenberg, W. P. Birmingham, B. Pardo, N. Hu, C. Meek, and G. Tzanetakis, "A comparative evaluation of search techniques for query- by-humming using the musart testbed," *J. of the American Soc. for Inform. Science and Technology*, vol. 58, no. 5, pp. 687–701, 2007.
- [2] A. Mesaros, T. Virtanen, and A. Klapuri, "Singer identification in poly- phonic music using vocal separation and pattern recognition methods," in *Proc. 8th Int. Conf. on Music Inform. Retrieval*, 2007, pp. 375–378.
- [3] J. L. Durrieu, G. Richard, and B. D. David, "An iterative approach to monaural musical mixture de-soloing," in *in Proc. IEEE Int. Conf. on Acoust., Speech and Signal Process. (ICASSP)*, Apr. 2009, pp. 105–108.
- [4] A. Ozerov, P. Philippe, F. Bimbot, and R. Gribonval, "Adaptation of bayesian models for single-channel source separation and its application to voice/music

- separation in popular songs,” *IEEE Trans. on Audio, Speech, and Language Process.*, vol. 15, no. 5, pp. 1564–1578, Jul 2007.
- [5] J. L. Durrieu, *Automatic transcription and separation of the main melody in polyphonic music signals*. Paris-Tech: Ph.D. dissertation, Telecom, 2010.
- [6] J. L. Durrieu, A. Ozerov, C. Fevotte, G. Richard, and B. D. David, “Main instrument separation from stereophonic audio signals using a source/filter model,” in *Proc. 17th European Signal Process. Conf. (EU-SIPCO)*, Glasgow, ago 2009, pp. 15–19.
- [7] J. Salamon and E. Gómez, “Melody extraction from polyphonic music signals using pitch contour characteristics,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, pp. 1759–1770, 08/2012 2012. [Online]. Available: <http://mtg.upf.edu/system/files/publications/SalamonGomezMelodyTASLP2012.pdf>
- [8] R. P. Paiva, T. Mendes, and A. Cardoso, “Melody detection in poly- phonic musical signals: Exploiting perceptual rules, note salience, and melodic smoothness,” *Computer Music J.*, vol. 30, pp. 80–98, 2006.
- [9] S. Koga, T. Inoue, and M. Fukumoto, “A Proposal for Intervention by User in Interactive Genetic Algorithm for Creation of Music Melody,” *2013 International Conference on Biometrics and Kansei Engineering*, pp. 129–132, Jul. 2013. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6603488>
- [10] Y. Journal, “A GENETIC ALGORITHM FOR COMPOSING MUSIC Dragan MATIĆ,” vol. 20, no. 1, pp. 157–177, 2010.
- [11] G. Reis, N. Fonseca, F. de Vega, and A. A. Ferreira, “Hybrid genetic algorithm based on gene fragment competition for polyphonic music transcription,” in *Proc. Conf. applications of evolutionary computing*, 2008, pp. 305–314.
- [12] B. Rafael, M. Affenzeller, and S. Wagner, “Application of an Island Model Genetic Algorithm for a Multi-track Music Segmentation Problem,” pp. 13–24.
- [13] J.-C. Hung and A.-C. Chang, “Combining genetic algorithm and iterative MUSIC searching DOA estimation for the CDMA system,” *Expert Systems with Applications*, vol. 38, no. 3, pp. 1895–1902, Mar. 2011. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0957417410007499>
- [14] A. D. Cheveigne and H. Kawahara, “Yin, a fundamental frequency estimator for speech and music,” *Journal of the Acoustical Society of America*, vol. 111, no. 4, pp. 1917–1930, 2002.
- [15] G. Kennedy and B. Kernfeld, “Aebersold, Jamey” in *B. Kernfeld. The new Grove dictionary of jazz*. New York: Grove’s Dictionaries Inc, 2002, vol. 1.
- [16] W. Marshall, *Best of Jazz*. Milwaukee, W, USA: Hall Leonard, 2000.
- [17] D. Bogdanov, N. Wack, E. Gómez, S. Gulati, P. Herrera, O. Mayor, G. Roma, J. Salamon, J. Zapata, and X. Serra, “Essentia: an audio analysis library for music information retrieval,” in *International Society for Music Information Retrieval Conference (ISMIR’13)*, Curitiba, Brazil, 04/11/2013 2013, pp. 493–498.
- [18] MATLAB, *version 7.12.0 (R2010a)*. Natick, Massachusetts: The MathWorks Inc., 2011.