

Recognizing Chords with EDS: Part One

Giordano Cabral¹, François Pachet², and Jean-Pierre Briot¹

¹ Laboratoire d'Informatique de Paris 6,
8 Rue du Capitaine Scott, 75015 Paris, France
{Giordano.CABRAL, Jean-Pierre.BRIOT}@lip6.fr

² Sony Computer Science Lab,
6 Rue Amyot, 75005 Paris, France
pachet@csl.sony.fr

Abstract. This paper presents a comparison between traditional and automatic approaches for the extraction of an audio descriptor to recognize chord into classes. The traditional approach requires signal processing (SP) skills, constraining it to be used only by expert users. The Extractor Discovery System (EDS) [1] is a recent approach, which can also be useful for non expert users, since it intends to discover such descriptors automatically. This work compares the results from a classic approach for chord recognition, namely the use of KNN-learners over Pitch Class Profiles (PCP), with the results from EDS when operated by a non SP expert.

1 Introduction

Audio descriptors express by mathematical formula a particular property of the sound, such as the tonality of a musical piece, the amount of energy in a given moment, or whether a song is instrumental or sung. Although the creation of each descriptor requires a different study, the design of a descriptor extractor normally follows the process of combining the relevant characteristics of acoustic signals (features) using machine learning algorithms. These features are often low-level descriptors (LLD), and the task usually requires important signal processing knowledge.

Since 2003, a heuristic-based approach became available through the Computer Science Lab of Sony in Paris, which developed the Extractor Discovery System (EDS). The system is based on genetic programming, and machine learning algorithms employed to automatically generate a descriptor from a database of sound files examples and their respective perceptive values. EDS can be used either by non experts or expert users. Non experts can use it as a tool to extract descriptors, even with minimal or no knowledge at all in signal processing. For example, movie makers have created classifiers of sound samples to be used in their films (explosions, car breaks, etc.). Experts can use the system to improve their results, starting from their solution and then controlling and guiding EDS. For instance, the perceived intensity of music titles can be more precisely revealed, taking as a starting point the mpeg7 audio features [2].

We are currently designing a guitar accompanier for “bossa nova” style. During the application development process, we ran into the problem of recognizing a chord, which turned out to be a good opportunity of comparing classical and EDS approaches. On the one hand, chord recognition is a well studied domain, with solid

results that can be considered as reference. On the other hand, current techniques use background knowledge that EDS (initially) does not have (pitches, harmony). Good EDS results would indicate the capacity of the system to deal with real world musical description cases.

We intend to compare the results from a standard technique of chord recognition (KNN learner over Pitch Class Profiles) and those from EDS, when operated by an inexperienced user (so called Naïve EDS) and by an expert user (so called Expert EDS). This paper presents the first part of this comparison, considering only the results obtained by the Naïve EDS. In the next section, we introduce the chord recognition problem. In section 3 we explain the most widely used technique. In section 4 we examine EDS, how it works and how to use it. Section 5 details the experiment. Section 6 shows and discuss the results. Finally, we draw some conclusions and point future works.

2 Chord Recognition

The ability of recognizing chords is important for many applications, such as interactive musical systems, content-based musical information retrieval (finding particular examples, or themes, in large audio databases), and educational software. Chord recognition means the transcription of a sound into a chord, which can be classified according to different levels of precision, from a simple distinction between maj and min chords to a complex set of chord types (maj, min, 7th, dim, aug, etc).

Many works can be mentioned here as the state of the art in chord recognition. [4] and [5] automatically transcribes chords from a CD recorded song. [3] deals with a similar problem: estimating the tonality of a piece (which is analogous to the maj/min). In most cases the same core technique is used (even if some variations may appear during the implementation phase): the computation of a Pitch Class Profile, or *chromagram*, and a subsequent machine learning algorithm to find patterns for each chord class. This technique has been applied to our problem, as we explain in next section.

3 Traditional Technique: Pitch Class Profiles

Most part of the works involving harmonic content (chord recognition, chord segmentation, tonality estimation) uses a feature called Pitch Class Profile (PCP) [6]. PCPs are vectors of low-level instantaneous features, representing the intensity of each pitch of the tonal scale mapped to a single octave. These vectors are calculated as follows: 1) a music recording is converted to a Fourier Transform representation (Fig1a to Fig1b). 2) the intensity of a pitch is calculated (Fig1b to Fig1d) by the magnitude of the spectral peaks, or by summing the magnitudes of all frequency bins that are located within the respective frequency band (Fig1c). 3) The equivalent pitches from different octaves are summed, producing a vector of 12 values (eventually 24 to deal with differences in tuning and/or to gain in performance), consequentially unifying various dispositions of a single chord class (Fig1e and Fig1f). For example, one can expect that the intensities of the frequencies corresponding to the notes C, E and G in the spectrum of a Cmaj would be greater than the others, independently on the particular voicing of the chord.




Fig. 1. Steps to compute a PCP. The signal is converted to Fast-Fourier representation; the FFT is divided into regions; the energy of each region is computed; the final vector is normalized.




Fig. 2. Example of the PCP for a Amaj7. Each column represents the intensity of a note, independently on the octave.




Fig. 3. Example of the PCP for a Cmaj7

The idea of using PCPs to chord recognition is that the PCPs of a chord follow a pattern, and that patterns can be learned from examples. Thus, machine learning (ML) techniques [9] can be used to generalize a classification model from a given database of labeled examples, in order to automatically classify new ones. So, for the PCP of a chord, the system will respond the most probable (or closest) chord class, given the examples previously learned. The original PCP implementation from Fujishima used a KNN learner [6], and more recent works [3] successfully used other machine learning algorithms.

4 EDS

EDS (Extractor Discovery System), developed at Sony CSL, is a heuristic-based generic approach for automatically extracting high-level music descriptors from acoustic signals. EDS is based on Genetic Programming [11], used to build extraction functions as compositions of basic mathematical and signal processing operators, such as *Log*, *Variance*, *FFT*, *HanningWindow*, etc. A specific composition of such operators is called feature (e.g. *Log (Variance (Min (FFT (Hanning (Signal))))*)), and a combination of features form a descriptor.

Given a database of audio signals with their associated perceptive values, EDS is capable to generalize a descriptor. Such descriptor is built by running a genetic search to find relevant signal processing features to match the description problem, and then machine learning algorithms to combine those features into a general descriptor model.




Fig. 4. EDS main interface

The genetic search performed by the system is intended to generate functions that may eventually be relevant to the problem. The best functions in a population are selected and iteratively transformed (by means of reproduction, i.e., constant variations, mutations, and/or cross-over), always respecting the pattern chosen by

the user. The default pattern is $!_x(\text{Signal})$, which means a function presenting any number of operations but a single value as result. The populations of functions keep reproducing until no improvement is achieved. At this point, the best functions are selected to be combined. This selection can be made both manually or automatically. For example, given a database of audio files labeled as ‘voice’/‘instrumental’, kept the default pattern, these are some possible functions that might be selected by the system:

```
Log10 (Range (Derivation (Sqrt (Blackman (MelBands (Signal, 24.0)))))
Square (Log10 (Mean (Min (Fft (Split (Signal, 4009)))))
```

Fig. 5. Some possible EDS features for characterizing a sound as vocal or instrumental

The final step in the extraction process is to choose and compute a model (linear regression, model trees, knn, locally weighted regression, neural networks, etc.) that combines all features. As an output, EDS creates an executable file, which classifies an audio file passed as argument.

In short, the user needs to 1) create the database, in which each recording is labeled as its correspondent class. 2) write a general pattern for the features and launch the genetic search. The pattern encapsulates the overall procedure of the feature. For example, $!_x(f:a(\text{Signal}))$ means that the signal is initially converted into the frequency domain, then some operation is applied to get a single value as a result. 3) select the appropriate features. 4) choose a model to combine the features. Although an expert user may drive the system (starting from an initial solution, including heuristics for the genetic search, etc), EDS has a fully automated mode, in which a default pattern is chosen, the most complementary features are selected and all models are computed. This mode is particularly attractive for non expert user, as he/she just needs to be able to create and label the database. That is the mode explored in this paper.

5 Bossa Nova Guitar Chords

Our final goal is to create a guitar accompanier in Brazilian “bossa nova” style. Consequently, our chord recognizer has examples of chords played with nylon guitar. The data was taken from D’accord Guitar Chord Database [10], a guitar midi based chord database. The purpose of using it was the richness of the symbolic information present (chord root, type, set of notes, position, fingers, etc.), which was very useful for labelling the data and validating the results. Each midi chord was rendered into a wav file using Timidity++ [12] and a free nylon guitar patch, and the EDS database was created according to the information found in D’accord Guitar database. Even though a midi-based database may lead to distortions in the results, we judge that the comparison between approaches is still valid.

5.1 Chord Classes

We tested the solutions with some different datasets, reflecting the variety of nuances that chord recognition may show:

AMaj/Min – classifies between major and minor chords, given a fixed root (La). There were 101 recordings, labelled in 2 classes.

Chord Type, fixed root – classifies among major, minor, seventh, minor seventh and diminished chords, given a fixed root (A or C). There were 262 samples, divided in 5 classes,

Chord Recognition – classifies major, minor, seventh, minor seventh and diminished chords, in any root. There were 1885 samples, labelled in 60 classes.

80% of each database is settled on as the training dataset and 20% as the testing dataset.

5.2 Pitch Class Profile

In our implementation of the pitch class profile, frequency to pitch mapping is achieved using the logarithmic characteristic of the equal temperament scale, as illustrated in Fig. 5. The intensity of each pitch is computed by summing the magnitude of all frequency bins that correspond to a particular pitch class. The same computation is applied to a white noise and the result is used to normalize the other PCPs.

$$Pitch = \frac{12 \times \log\left(\frac{|H_{z_{bin}}|}{440}\right)}{\log(2)}$$

Fig. 6. Frequency to pitch mapping

For the *chord recognition* database, PCPs were rotated, meaning that each PCP was computed 12 times, one time for each possible rotation (for instance, a Bm is equivalent to a Am rotated twice). After the PCP computation, several machine learning algorithms could have been applied. We implemented 2 simple solutions. The first




Fig. 7. Example of a template PCP for a C chord class

one calculates a default (template) PCP to each chord class. Then, the PCP of a new example can be matched up to the template PCP, and the most similar one is retrieved as the chord.

The second one uses the k-nearest neighbours algorithm (KNN), with maximum of 3 neighbours. KNNs have been used since the original PCP implementation and have proved to be at least one of the best learning algorithms for this case [3].

5.3 EDS

The same databases were loaded in EDS. We ran a fully automated extraction, keeping all default values. The system generated the descriptor without any help from the user, obtaining the results we call EDS Naïve, because they correspond to the results that a naïve user would achieve.

6 Results and Discussion

The results achieved by us are presented in the table 1. Rows represent the different databases. Columns represent the different learning techniques. The percent values indicate the number of correctly classified instances over the total number of examples in the testing database.

As we can see, EDS gets really close to classical approaches when the root is known, but disappoints when the whole problem is presented. It seems that a combination of low level functions is capable of finding different patterns in the same root, but the current palette of signal processing functions in EDS is not sufficient to generalize harmonic information. Sections 6.1, 6.2 and 6.3 detail the features that were found.

Table 1. Percentage of correctly classified instances for the different databases using the studied approaches

Approach Database	PCP Template	KNN	EDS
Maj/Min (fixed root)	100%	100%	90.91%
Chord Type (fixed root)	89%	90.62%	87.5%
Chord Recognition	53.85%	63.93%	40.31%

6.1 Case 1: Major/Minor Classifier, Fixed Root

Figure 5 shows the selected features for the *Amaj/min* database. The best model obtained was a KNN of 1 nearest neighbour, equally weighted, absolute error (see [9] for details). The descriptor reached 90.91% of the performance of the best traditional classifier.

```

EDS1: Power (Log10 (Abs (Range (Integration (Square (Mean (FilterBank
(Normalize (Signal), 5.0)))))), -1.0)

EDS2: Power (Log10 (Abs (Range (Sqrt (Bartlett (Mean (FilterBank
(Normalize (Signal), 9.0)))))), -1.0)

EDS3: Sqrt (Range (Integration (Hanning (Square (Mean (Split (Signal,
3736.0))))))

EDS4: Arcsin (Sqrt (Range (Integration (Mean (Split (Normalize (Sig-
nal), 5862.0))))))

EDS5: Log10 (Variance (Integration (Bartlett (Mean (FilterBank (Nor-
malize (Signal), 5.0))))))

EDS6: Power (Log10 (Abs (Range (Integration (Square (Sum (FilterBank
(Normalize (Signal), 9.0)))))), -1.0)

EDS7: Square (Log10 (Abs (Mean (Normalize (Integration (Normalize
(Signal))))))

EDS8: Arcsin (Sqrt (Range (Integration (Mean (Split (Normalize (Sig-
nal), 8913.0))))))

EDS9: Power (Log10 (Abs (Range (Sqrt (Bartlett (Mean (FilterBank
(Normalize (Signal), 3.0)))))), -1.0)

```

Fig. 8. Selected features for the *Amaj/min* chord recognizer

6.2 Case 2: Chord Type Recognition, Fixed Root

Figure 6 shows the selected features for the chord type database. The best model obtained was a GMM of 14 gaussians and 500 iterations (see [9] for details). The descriptor reached 96,56% of the performance of the best traditional classifier.

```

EDS1: Log10 (Abs (RHF (Sqrt (Integration (Integration (Normalize
(Signal))))))

EDS2: Mean (Sum (SplitOverlap (Sum (Bartlett (Split (Signal,
1394.0))), 4451.0, 0.5379660839449434))

EDS3: Power (Log10 (Abs (RHF (Normalize (Integration (Integration
(Normalize (Signal)))))), 6.0)

EDS4: Power (Log10 (RHF (Signal)), 3.0)

EDS5: Power (Mean (Sum (SplitOverlap (Sum (Bartlett (Split (Signal,
4451.0))), 4451.0, 0.5379660839449434))), 3.0)

```

Fig. 9. Selected features for the Chord Type recognizer

6.3 Case 3: Chord Recognition

Figure 7 shows some of the selected features for the chord recognition database. The best model obtained was a KNN of 4 nearest neighbours, weighted by the inverse of the distance (see [9] for details). The descriptor reached 63,05% of the performance of the best traditional classifier. It is important to notice that 40,31 % is not necessarily a

bad result, since we have 60 possible classes. In fact, 27,63% of the wrongly classified instances were due to mistakes between relative majors and minors (e.g. C and Am); 40,78% due to other usual mistakes (e.g. C and C7; C° and Eb°; C and G); only 31,57% were caused by unexpected mistakes. Despite these remarks, the comparative results are significantly worse than the previous ones.

```
EDS1: Square (Log10 (Abs (Sum (SpectralFlatness (Integration (Split
(Signal, 291.0)))))))

EDS4: Power (Log10 (Abs (Iqr (SpectralFlatness (Integration (Split
(Signal, 424.0)))))), -1.0)

EDS9: Sum (SpectralRolloff (Integration (Hamming (Split (Signal,
4525.0))))))

EDS10: Power (Log10 (Abs (Median (SpectralFlatness (Integration
(SplitOverlap (Signal, 5638.0, 0.7366433546185794)))))), -1.0)

EDS12: Log10 (Sum (MelBands (Normalize (Signal)), 7.0))

EDS13: Power (Median (Normalize (Signal)), 5.0)

EDS14: Rms (Range (Hann (Split (Signal, 9336.0))))

EDS15: Power (Median (Median (Split (Sqrt (Iqr (Hamming (Split (Sig-
nal, 2558.0))))), 4352.0))), 1.5)

EDS17: Power (HFC (Power (Correlation (Normalize (Signal), Signal),
4.0)), -2.0)

EDS18: Square (Log10 (Variance (Square (Range (Mfcc (Square (Hamming
(Split (Signal, 9415.0))), 2.0))))))

EDS19: Variance (Abs (Median (Hann (FilterBank (Peaks (Normalize
(Signal)), 5.0))))))

EDS21: MaxPos (Sqrt (Normalize (Signal)))

EDS22: Power (Log10 (Abs (Iqr (SpectralFlatness (Integration (Split
(Signal, 4542.0)))))), -1.0)
```

Fig. 10. Some of the selected features for the chord recognizer

6.4 Other Cases

We also compared the three approaches on other databases, as we can see in the table 2. *MajMinA* is the major/minor classifier, root fixed to A. *ChordA* is the chord type recognizer, root fixed to A. *ChordC* is the chord type recognizer, root fixed to C. *RealChordC* is the same chord type recognizer in C, but the testing dataset is composed by real audio recordings (samples of less than 1 second of chords played in a nylon guitar), instead of midi rendered audio. Curiously, in this case, the EDS solution worked better than the traditional one (probably due to an alteration in tuning in the recorded audio). *Chord* is the chord recognition database. *SmallChord* is a smaller dataset (300 examples) for the same problem. Notice that in this case EDS outperformed KNN and PCP Template. In fact, the EDS solution does not improve very much when passing from 300 to 1885 examples (from 38,64% to 40,31%), while the

KNN solution goes from 44% to 63,93%. Finally, *RealChord* has the same training set from the *Chord* database, but is tested with real recorded audio.

The results from these databases confirm the trend of the previous scenario. The reading of the results indicates that the effectiveness of the EDS fully automated descriptor extraction depends on the domain it is applied to. Even admitting that EDS (in its current state) is only partially suited to non expert users, we must take into account that EDS currently uses a limited palette of signal processing functions, which is being progressively enhanced. Since EDS didn't have any information about tonal harmony, it was already expected that it would not reach the best results. Even though, the results obtained by the chord recognizer with a fixed root show the power of the tool.

Table 2. Comparison between the performance of the EDS and the best traditional classifier for a larger group of databases. Comparative performance = EDS performance / traditional technique performance.

DB NAME	Comparative Performance
MajMinA	90,91%
ChordA	94,38%
ChordC	96,56%
Chord	63,05%
SmallChord	87,82%
RealChordC	116,66%
RealChord	55,16%

7 Conclusion and Future Works

In this paper we compared the performance of a standard chord recognition technique and the EDS approach. The chord recognition was specifically related to nylon guitar samples, since we intend to apply the solution to a Brazilian style guitar accompanier. The standard technique was the Pitch Class Profiles, in which frequency intensities are mapped to the twelve semitone pitch classes, and then uses KNN classification to chord templates. EDS is an automatic descriptor extractor system that can be employed even if the user does not have knowledge about signal processing. It was operated in a completely naïve way so that the solution and the results would be similar to those obtained by a non expert user.

The statistical results reveal a slight deficit of EDS for a fixed root, and a greater gap when the root is not known a priori, showing its dependency on primary operators. An initial improvement is logically the increase of the palette of functions. Currently, we are implementing tonal harmony operators such as chroma and pitchBands, which we suppose will provide much better results. Additionally, as the genetic search in EDS is indeed an optimisation algorithm, if the user starts from a

good solution, it will be expected that the algorithm makes it even better. The user can also guide the function generation process, via more specific patterns and heuristics.

With these actions, we intend to perform the second part of the comparison we started in this paper – between the traditional techniques and EDS operated by a signal processing expert.

Acknowledgements

We would like to thank all the team at Sony CSL Paris, particularly Anthony Beurivé, Jean-Julien Aucouturier and Aymeric Zils for their support and assistance with EDS; and a special thanks to Tristan Jehan for his help in the conception and implementation of the algorithms.

References

1. Pachet, F. and Zils, A. "Automatic Extraction of Music Descriptors from Acoustic Signals", Proceedings of Fifth International Conference on Music Information Retrieval (ISMIR04), Barcelona, 2004.
2. Zils, A. & Pachet, F. "Extracting Automatically the Perceived Intensity of Music Titles", Proceedings of the 6th COST-G6 Conference on Digital Audio Effects (DAFX03), 2003.
3. Gómez, E. and Herrera, P. "Estimating the tonality of polyphonic audio files: cognitive versus machine learning modelling strategies", Proceedings of the 5th International Conference on Music Information Retrieval (ISMIR04), Barcelona, 2004.
4. Sheh, A. and Ellis, D. "Chord Segmentation and Recognition using EM-Trained Hidden Markov Models", Proceedings of the 4th International Symposium on Music Information Retrieval (ISMIR03), Baltimore, USA, 2003.
5. Yoshioka, T., Kitahara, T., Komatani, K., Ogata, T. and Okuno, H. "Automatic chord transcription with concurrent recognition of chord symbols and boundaries", Proceedings of the 5th International Conference on Music Information Retrieval (ISMIR04), Barcelona, 2004.
6. Fujishima, T. "Real-time chord recognition of musical sound: a system using Common Lisp Music", Proceedings of International Computer Music Conference (ICMC99), Beijing, 1999.
7. Bartsch, M. A. and Wakefield, G. H. "To catch a chorus: Using chromabased representation for audio thumbnailing", Proceedings of International Workshop on Applications of Signal Processing to Audio and Acoustics, Mohonk, USA, 2001.
8. Pardo, B., Birmingham, W. P. "The Chordal Analysis of Tonal Music", The University of Michigan, Department of Electrical Engineering and Computer Science Technical Report CSE-TR-439-01, 2001.
9. Mitchell, T. "Machine Learning", The McGraw-Hill Companies, Inc. 1997.
10. Cabral, G., Zanforlin, I., Santana, H., Lima, R., & Ramalho, G. "D'accord Guitar: An Innovative Guitar Performance System", in Proceedings of Journées d'Informatique Musicale (JIM01), Bourges, 2001.
11. Koza, J. R. "Genetic Programming: on the programming of computers by means of natural selection", Cambridge, USA, The MIT Press.
12. Gómez, E. Herrera, P. "Automatic Extraction of Tonal Metadata from Polyphonic Audio Recordings", Proceedings of 25th International AES Conference, London, 2004.
13. Website: <http://timidity.sourceforge.net/>