

# MIREX SUBMISSIONS FOR AUDIO CHORD DETECTION (NO TRAINING) AND STRUCTURAL SEGMENTATION

Matthias Mauch, Katy Noland, Simon Dixon

Queen Mary University of London, Centre for Digital Music

{matthias.mauch, kathy.noland, simon.dixon}@elec.qmul.ac.uk

## ABSTRACT

This paper describes our approach to chord extraction from audio, a variant of which was submitted to the 2009 MIREX Chord Detection Task (No Training), and achieved the top ranking of 71.2%. The structural segmentation algorithm is a pre-processing step for the chord extraction, and was also submitted separately for the Structural Segmentation Task. It also achieved the top ranking in that category with a pairwise F-measure of 0.6.

Chord extraction from audio is a well-established music computing task, and many valid approaches have been presented in recent years that use different chord templates, smoothing techniques and musical context models. The present work exploits the repetitive structure of songs to enhance chord extraction, by combining chroma information from multiple occurrences of the same segment type. Our results published in [1] show that the method produces consistent and more readily readable chord labels and provides a statistically significant boost in label accuracy.

## 1. INTRODUCTION

The automatic extraction of chords from audio has applications in music retrieval, cognitive musicology, and automatic generation of lead sheets. In this work we present a technique that allows us to generate more authentic lead sheets than previously possible with automatic methods, by making use of musical structure. Much of musical structure is defined by repetition, a core principle in music [2, p. 229].

In popular songs a repeated *verse-chorus* format is common, in which the chord sequence is the same in all sections of the same type. In lead sheets, for better readability these sections would normally only be notated once, with repeats indicated. Our method mirrors this improvement by assigning the same chord progression to repeated sections. In addition, having found repeating sections, we have available several instances of a given chord sequence from which to estimate the chords, which can lead to an improvement in estimation accuracy.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

© 2009 International Society for Music Information Retrieval.

In Section 2 we describe related work. In Section 3 we describe the chord extraction method used and present a new segmentation technique that is tailored to our task of finding repeated chord sequences. In Section 4 we discuss some more general issues. For examples and results, please refer to our ISMIR paper [1].

## 2. RELATED WORK

The majority of approaches to automatic chord estimation rely on frame-wise chroma features [3] as a representation of the relative energy in each pitch class for a given time window, then apply some further processing to estimate the chords. When template-matching is used to identify chords, additional smoothing over time, for example by a median filter [4], is necessary due to musical variation and noise. Inference in hidden Markov models (HMMs) [5] simultaneously performs template-matching and smoothing. These methods treat chords as isolated features of the music, which is a considerable simplification. In reality, chords are heard in context, together with the melody, key, rhythm, form, instrumentation, and other attributes. Some chord estimation methods account for additional musical attributes during the estimation process such as key [6], bass [7], key and rhythm [8, 9], or even key, rhythm and bass together [10], which is a step towards a unified music analysis model.

In this work we extend the concept of unified music analysis by using repetition in the structure to enhance chord estimation. Dannenberg [11] shows that knowledge of the musical structure can greatly improve beat tracking performance, but to our knowledge the principle has not yet been applied to chord estimation.

Previous automatic music structure extraction techniques include those that primarily search for section boundaries, indicated by a sudden change in the feature of interest, which could be timbre [12], spectral evolution [13], or combinations of features [14]. A common approach is to cluster together frames that are similar, then label contiguous similar frames as a segment. However, this relies on a particular feature remaining approximately constant for the duration of a section. We are interested in chords, which do change during a section, so an approach that searches for repeated progressions [15, 16] is more appropriate for our purposes. Methods using this paradigm rely on a self-similarity matrix [17], which is a symmetric, square matrix that contains a measure of the similarity

between every pair of frames. Repeated sections appear as parallel diagonal lines, and can be extracted with some post-processing, such as application of a low pass filter to reduce noise [18] followed by a thresholding operation to find contiguous frames with high similarity. In Section 3.3 we present a new variation which is similar to algorithms proposed by Ong [19] and Rhodes and Casey [20] and extracts repeated segments of equal length.

### 3. METHOD

In a song, we call a chord sequence that describes a section such as the verse or chorus a *segment type*. Any segment type may occur one or more times in a song and we call each occurrence a *segment instance*. To make use of segment repetition as part of the chord estimation process, we rely on segment types whose instances are not only harmonically very similar, but also have the same length in beats (see Section 3.4). This is not required of a general purpose segmentation task, and hence generic segmentations are not directly utilisable. In Section 3.2 we describe how we preprocess manual segmentations to meet our needs. For automatic segmentation we choose to implement our own algorithm, which fulfills the above requirements by design (Section 3.3). First, we describe the method for calculating our basic features, beat-synchronous chromagrams (Section 3.1).

#### 3.1 Beat-Synchronous Chromagrams

The automatic segmentation and chord estimation algorithms both rely on chroma features that are synchronised to the musical beat. The features represent the importance of each pitch class at the current beat. The initial, short chroma frames are generated from a note salience representation similar to a constant-Q transform, at a hopsize of 512 samples (46 ms) from audio that has been down-sampled to 11025 Hz. For the chord extraction algorithm we split the salience representation to obtain separate bass and treble chromagrams, but the chromagram used by the segmentation algorithm covers both the bass and the treble range. For details see [21].

In order to produce beat-synchronous chromagrams we obtain a single chroma vector for each beat by taking the median (in the time direction) over all the chroma frames falling between two consecutive beat times. We use one of two sorts of beat times: manual or automatic. The collection of manual beat annotations [22] covers 125 songs performed by the rock group The Beatles. The automatic beat times were extracted using Davies’s automatic beat-tracker [23] on the same set of songs.

#### 3.2 Manual Structural Segmentation

The manual structural segmentations [22] cover the same 125 songs by The Beatles as we have beat annotations for: 29 songs were annotated for a previous project<sup>1</sup>, and 96

<sup>1</sup>Segmentations available at <http://www.elec.qmul.ac.uk/digitalmusic/downloads/index.html#segment>.

were newly annotated for the present work. The basis for all annotations are Pollack’s song analyses [24].

Every song contains several segment types, some of which have multiple instances. In some songs, the instances of a segment type differ in length. In that case, to fulfill the requirement of equal length instances, the segment type is divided to create one or more new segment types whose instances all have the same length. This may result in new segment types having only one instance in the song.

#### 3.3 Automatic Segmentation Algorithm

The automatic segmentation method has two main steps: finding approximately repeated chroma sequences in a song, and a greedy algorithm to decide which of the sequences are indeed segments. We calculate the Pearson correlation coefficients between every pair of chroma vectors, which together represent a beat-wise self-similarity matrix  $R = (r_{ij})$  of the whole song. This is similar to the matrix of cosine distances used by Ong [19]. In the similarity matrix, parallel diagonal lines indicate repeated sections of a song. In order to eliminate short term noise or deviations we run a median filter of length 5 (typically just more than one bar) diagonally over the similarity matrix. This step ensures that *locally* some deviation is tolerated.

We perform a search of repetitions over all diagonals in the matrix over a range of lengths. We assume a minimum length of  $m_1 = 12$  beats and a maximum length of  $m_M = 128$  beats for a segment, leading to a very large search space. We minimise the number of elements we have to compare by considering as section beginnings only those beats that have a correlation  $r$  greater than a threshold  $t_r$ , and assuming that section durations are quantised to multiples of four beats. We found that a value of  $t_r = 0.65$  worked well. In future work we would like to learn  $t_r$  from data. We further reduce the search space by allowing segments to start only at likely bar beginnings. Likely bar beginnings are beats where the convolution of a function representing the likelihood of a change in harmony, and a kernel with spikes every two beats has a local maximum (details in [21]).

To assess the similarity of a segment of length  $l$  starting at beat  $i$  to another one of the same length starting at  $j$  we consider the diagonal elements

$$D_{i,j,l} = (r_{i,j}, r_{i+1,j+1}, \dots, r_{i+l,j+l}) \quad (1)$$

of the matrix  $\mathcal{R}$ . If the segments starting at  $i$  and  $j$  are exactly the same, then  $D_{i,j}$  will be a vector of ones, and hence we can characterise a perfect match by

$$\min\{D_{i,j,l}\} = 1. \quad (2)$$

To accommodate variation arising in a practical situation, we relax the requirement (2) by using the empirical  $p$ -quantile function<sup>2</sup> instead of the minimum (which is the 0-quantile), and choosing a segment threshold  $t_s$  lower than

<sup>2</sup><http://www.mathworks.com/access/helpdesk/help/toolbox/stats/quantile.html>

unity. The triple  $(i, j, l)$  hence describes a repetition, if

$$\text{quantile}_p\{D_{i,j,l}\} > t_s. \quad (3)$$

The two parameters  $p = 0.1$  and  $t_s = 0.6$  are chosen empirically. In future work we would like to learn these values from the ground truth data. The set of repetitions  $\mathcal{R}_{il} = \{j : \text{quantile}_p\{D_{i,j,l}\} > t_s\}$  is then added to a list  $\mathcal{L}$  of repetition sets, if it has more than one element  $j$ , i.e. if it actually describes at least one repetition. If two segments  $(i, j_1, l)$  and  $(i, j_2, l)$  overlap, only the index of the one with the higher score is retained in  $\mathcal{R}_{il}$ .

Each of the sets  $\mathcal{R}_{il}$  represent a potential segment type, and its elements represent the start beats of instances of that segment type. However, there are typically many more repetition sets than there are segment types. To find repetition sets relating to actual segment types we use the heuristic of a music editor who tries to save paper: he will first take the repetition set in which  $l \times |\mathcal{R}_{il}|$  is maximal, and then repeat this kind of choice on the remaining segments of the song, resulting in a greedy algorithm. The only exception to that rule is the case in which he finds that a sub-segment of a repetition is repeated more often than the whole segment. He then chooses the  $\mathcal{R}_{il}$  pertaining to the sub-segment.

### 3.4 Using Repetition Cues in Chord Extraction

We use structural segmentation to combine several instances of a segment type in a song and then infer a single chord sequence from the combination.

The baseline is an existing chord labelling method [10], which extracts chords from beat-synchronous treble and bass chromagrams. Using a dynamic Bayesian network [25] similar to a hierarchical hidden Markov model the network jointly models metric position, chords and bass pitch class and infers the most probable sequence from the beat-synchronous chromagrams of the whole song. The method models four different chord classes: major, minor, diminished and dominant<sup>3</sup>.

In order to integrate the knowledge of repeating segments, we split the chromagram for the whole song into smaller chromagram chunks, each belonging to one segment instance. If a segment type has more than one instance, all its chromagram chunks are averaged by taking the mean of the respective elements, thus creating a new chromagram chunk representing all instances of the segment type. The chord extraction is then performed on the newly generated chromagram chunk, and the estimated chords are transcribed as if they had been extracted at the individual segment instances.

## 4. DISCUSSION

The method presented here is not tied to the individual algorithms. Using other chord extraction or segmentation methods could further improve results and shed more light on the performance of its constituent parts. As mentioned

in Section 3.3 we plan to investigate the effects of training some of the segmentation parameters. It would also be interesting to determine whether using the median (instead of the mean) to average chromagram chunks would lead to improvements.

The present work focussed on early rock music. We expect that—given a good segmentation—improvements in recognition results could be even greater for jazz: while the extraction of chords in jazz is more difficult than in rock music due to improvisation and more complex chord types, the repetition of segment types is often more rigid.

The method to share information globally between segments we used for this work is a simple one. Integrating this process with the chord extraction itself would be a more elegant solution.

## 5. REFERENCES

- [1] Matthias Mauch, Katy C. Noland, and Simon Dixon. Using musical structure to enhance automatic chord transcription. In *Proceedings of the 10th International Conference on Music Information Retrieval, Kobe, Japan, 2009*.
- [2] David Huron. *Sweet Anticipation: Music and the Psychology of Expectation*. MIT Press, 2006.
- [3] Takuya Fujishima. Real time chord recognition of musical sound: a system using Common Lisp Music. In *Proceedings of the International Computer Music Conference (ICMC)*, pages 464–467, 1999.
- [4] Christopher Harte and Mark Sandler. Automatic chord identification using a quantised chromagram. In *Proceedings of 118th Convention*. Audio Engineering Society, 2005.
- [5] Juan P. Bello and Jeremy Pickens. A Robust Mid-level Representation for Harmonic Content in Music Signals. In *Proceedings of the 6th International Conference on Music Information Retrieval, ISMIR 2005, London, UK*, pages 304–311, 2005.
- [6] Kyogu Lee and Malcolm Slaney. Acoustic Chord Transcription and Key Extraction From Audio Using Key-Dependent HMMs Trained on Synthesized Audio. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(2):291–301, February 2008.
- [7] Matthias Mauch and Simon Dixon. A Discrete Mixture Model for Chord Labelling. In *Proceedings of the 9th International Conference on Music Information Retrieval, ISMIR 2008, Philadelphia, USA*, pages 45–50, 2008.
- [8] Arun Shenoy and Ye Wang. Key, chord, and rhythm tracking of popular music recordings. *Computer Music Journal*, 29(3):75–86, September 2005.
- [9] Hélène Papadopoulou and Geoffroy Peeters. Simultaneous estimation of chord progression and downbeats from an audio file. In *Proceedings of the 2008 ICASSP Conference*, pages 121–124, 2008.

<sup>3</sup> strictly speaking: major with a minor seventh

- [10] Matthias Mauch and Simon Dixon. Simultaneous estimation of chords and musical context from audio. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(2), February 2010.
- [11] Roger B. Dannenberg. Toward automated holistic beat tracking, music analysis, and understanding. In *Proceedings of the 6th International Conference on Music Information Retrieval*, London, 2005.
- [12] S. Abdallah, K. Noland, M. Sandler, M. Casey, and C. Rhodes. Theory and evaluation of a Bayesian music structure extractor. In *Proceedings of the 6th International Conference on Music Information Retrieval, ISMIR 2005, London, UK*, pages 420–425, 2005.
- [13] G. Peeters, A. La Burthe, and Xavier Rodet. Toward automatic music audio summary generation from signal analysis. In *Proceedings of the 3rd International Conference on Music Information Retrieval*, Paris, 2002.
- [14] Namunu C. Maddage. Automatic structure detection for popular music. *IEEE Multimedia*, 13(1):65–77, 2006.
- [15] Meinard Müller and Frank Kurth. Towards structural analysis of audio recordings in the presence of musical variations. *EURASIP Journal on Advances in Signal Processing*, 2007.
- [16] Masataka Goto. A chorus-section detecting method for musical audio signals. In *Proceedings of the 2003 IEEE Conference on Acoustics, Speech and Signal Processing*, pages 437–440, 2003.
- [17] Jonathan Foote. Visualizing music and audio using self-similarity. In *Proceedings of the 7th ACM International Conference on Multimedia (Part 1)*, pages 77–80, 1999.
- [18] Mark A. Bartsch and Gregory H. Wakefield. Audio thumbnailing of popular music using chroma-based representations. *IEEE Transactions on Multimedia*, 7(4), February 2005.
- [19] Bee Suan Ong. *Structural Analysis and Segmentation of Music Signals*. PhD thesis, Universitat Pompeu Fabra, 2006.
- [20] Christophe Rhodes and Michael Casey. Algorithms for determining and labelling approximate hierarchical self-similarity. In *Proceedings of the 2007 ISMIR Conference, Vienna, Austria*, pages 41–46, 2007.
- [21] Matthias Mauch. A chroma extraction method and a harmonic change detection function. Technical report, Queen Mary, University of London, 2009. Available at <http://www.elec.qmul.ac.uk/digitalmusic/papers/2009/Mauch09-C4DM-TR-09-05.pdf>.
- [22] M. Mauch, C. Cannam, M. Davies, S. Dixon, C. Harte, S. Kolozali, D. Tidhar, and M. Sandler. Omras2 metadata project 2009. In *Late-breaking session at the 10th International Conference on Music Information Retrieval, Kobe, Japan, 2009*.
- [23] Matthew Davies. *Towards Automatic Rhythmic Accompaniment*. PhD thesis, Queen Mary University of London, London, UK, August 2007.
- [24] Alan W. Pollack. Notes on... series, 1995. Available at <http://www.recmusicbeatles.com>.
- [25] Kevin P Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, University of California, Berkeley, 2002.