DISCOVERING CHORD IDIOMS THROUGH BEATLES AND REAL BOOK SONGS

Matthias Mauch Simon Dixon Christopher Harte
Queen Mary, University of London,
Centre for Digital Music
matthias.mauch@elec.qmul.ac.uk

Michael Casey Benjamin Fields
Goldsmiths, University of London,
Computing Department

ABSTRACT

Modern collections of symbolic and audio music content provide unprecedented possibilities for musicological research, but traditional qualitative evaluation methods cannot realistically cope with such amounts of data. We are interested in harmonic analysis and propose key-independent *chord idioms* derived from a bottom-up analysis of musical data as a new subject of musicological interest. In order to motivate future research on audio chord idioms and on probabilistic models of harmony we perform a quantitative study of chord progressions in two popular music collections. In particular, we extract common subsequences of chord classes from symbolic data, independent of key and context, and order them by frequency of occurrence, thus enabling us to identify chord idioms. We make musicological observations on selected chord idioms from the collections.

1 INTRODUCTION

Traditional musicology consists of qualitative studies using small data sets, so that it is not possible to ascertain whether the conclusions drawn from the study are representative of a broader corpus of music. Harmonic analysis is no exception to that rule, and there is plenty of literature on harmony in Western music and Jazz (e.g. [3]). One of the disadvantages of this approach is that the choice of data and its interpretation are subjective. Music Information Retrieval methods provide us with increasingly powerful tools that can be applied to strip some of such subjectivity from the analyses by quantitatively evaluating features over large collections of music. Only in recent years has considerable effort been put into the automatic analysis of chord changes in audio and symbolic representations. Most of the efforts in audio chord analysis are concerned with the actual extraction of chords (e.g. [5], [2] and [8]) – which is an interesting and very difficult task – but do not address musicological questions.

In the symbolic domain too, efforts have mainly been directed towards chord transcription. Examples include MIDI-driven harmony retrieval such as the Melisma Harmony Program [10] or harmonic labelling with Bayesian Model Selection [9]. On the other hand, Pachet [7] considers hand-annotated chord labels and infers a notion of surprise in chord sequences and a set of chord substitution rules.

2 THE STUDY

We follow Pachet by using manually extracted chord labels for our analyses. We deliberately avoid any analysis of tonality or other high-level features because we believe that they are coded implicitly in chord sequences of sufficient length. That is, the chord sequences themselves represent the evolution of harmony over time.

Our aim is to create an inventory (or: dictionary) of chord sequences together with an analysis of their statistical frequencies in order to discover *chord idioms*, prominent chord sequences in a particular style, genre or historical period. This inventory will also function as a basis for further probabilistic research into harmonic structures.

We examine two collections of manually labelled chord symbols in text form. One is Harte's Beatles Chord Database [6] consisting of all the chords of all 180 songs featured on original Beatles studio albums. The database includes start and end times of chords in songs relative to the original Beatles' recordings. The other collection is a transcription of the chords of 244 Jazz standards from the Real Book [12].

2.1 Chord Classes

Although the notations in the two collections are relatively compatible, the first step was to translate the chord labels into a standardised format. A chord label consists of a root note $r \in \{0, \ldots, 11\}$, and a chord type represented as a tuple (b, c), where $b \in \{0, \ldots, 11\}$ describes the bass note in semitones above the chord root, and $c \in \{C_1, \ldots, C_{N_C}\}$ represents the remaining chord structure. More precisely, each C_i is the set of pitches in the chord, expressed in semitones above the root. Thus a chord type identifies the structure of the chord, e.g. a minor triad in root position $(0, \{0, 3, 7\})$, or the second inversion of a dominant 7th chord $(7, \{0, 4, 7, 10\})$.

Although studying the use of specific chord types is in itself interesting, for discovering general chord idioms it is more advantageous to create categories of chords that have a similar harmonic function (see also section 3.3). In order to do so we establish sets of chord classes, which we then use instead of the actual chord types.

one class	C_1	all chords (only root note differences count)
five classes	C_1	major chords, C5 and C2
	C_2	minor chords
	C_3	half-diminished and diminished chords
	C_4	augmented chords, incl. "7+" etc.
	C_5	chords with a suspended 4th
full	C_1	$\{0,4,7\}$ (major)
	C_2	$\{0,3,7\}$ (minor)
	C_3	$\{0,4,7,10\}$ (7th)
	C_4	$\{0,4,7,10,14\}$ (9th)
	...	
	C_{N_C}	$\{\dots\}$

Table 1. Three sets of chord classes, representing three different levels of generality.

From the many possible sets of chord classes, we propose the musically plausible sets in table 1. For example, using the *five classes* set the chord types in the sequence

```
Dm7     | G13     | CMaj7/E     | A9     |
```

from *Afternoon in Paris* would be translated to $(0, C_2)$, $(0, C_1)$, $(4, C_1)$ and $(0, C_1)$. In the following we will always refer to the *five classes* set, except where stated.

2.2 Chord Changes and Sequences

Let $s = (b_i, c_i)_{i=1\dots n}$ be a sequence of n successive chord types, and $d = (d_1, \dots, d_{n-1}),\ d_i \in \{0, \dots, 11\}$ the root differences (modulo 12) between them. Then we call $S = (s, d)$ a *length n chord sequence*. For instance, a simple chord change is represented by a length 2 chord sequence. With the difference in semitones between `D` and `G` being 5, the first chord change in the above example is

$$s = ((0, C_2), (0, C_1)), \quad d_1 = 5,$$

or, in a more accessible way,

$$\text{minor} \xrightarrow{+5} \text{major}.$$

This representation has the advantage of not containing the actual chord roots, so it is invariant to transposition.

2.3 Chord idioms

Not all chord sequences of same length are equally important. The expression *chord idioms* denotes chord sequences that we deem special. For instance, we consider the ii-V-I progression one central idiom in the Jazz literature. According to [1] an idiom in natural language is

> "an expression in the usage of a language that is peculiar to itself [...] in having a meaning that cannot be derived from the conjoined meanings of its elements [...]"
> "a style or form of artistic expression that is characteristic of [...] a period [...] "

Approaching harmony in an analogous way, we take a bottom-up approach in order to investigate harmonic usage, rather than a top-down approach based on some (unknown) underlying rule model.

Our conception of chord idioms guides the choice of the length of the chord sequences. For instance, a length 2 chord sequence (i.e. a chord change) is not appropriate because it is too short to capture local tonality and is akin to a word in spoken language rather than to an idiom. On the other hand, chord sequences of length greater than 8 in small or medium size collections will become very rare, which prevents them from being "characteristic" of a collection (as required in the definition) and makes a statistical analysis nearly impossible. The algorithm itself does not require any particular sequence length, but for the present article we (somewhat arbitrarily) chose a length of 4 chords in order to illustrate our methodology. Four is also a typical phrase length. In section 2.4 we introduce two measures for chord sequence frequency whose combination might help us to recognise chord idioms automatically.

2.4 Estimator for Chord Sequence Probability

It is not straight-forward to find a good estimator for the probability of a chord sequence as the overall relative frequency is biased towards repetitive patterns. Instead, for every chord sequence we estimate

1. the probability p_i of a chord sequence S_i occurring at least once in a song.
2. the probability q_i that a chord sequence in a song is S_i, conditional on S_i occurring at least once in the song.

These probabilities are based on the frequency at which the chord sequences occur in a given music database. Note however, that – due to the fact that we consider overlapping chord sequences – only p_i is robust to alternating chord changes, while q_i will count the chord sequence

$$\text{maj} \xrightarrow{+5} \text{maj} \xrightarrow{+7} \text{maj} \xrightarrow{+5} \text{maj}$$

three times in the following chord progression, not just twice!

```
D   G   |D   G   |D   G   |D   G   |
```

3 RESULTS

3.1 Chord Class Frequencies

Figure 1 shows the distribution of chord classes for the two data sets. It is evident that major chords prevail in

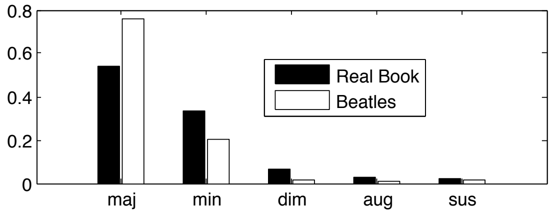

Figure 1. Frequencies for the *five classes* in the Real Book and the Beatles Chord Database

both collections. In the Beatles' songs 76% of all chords are classified major and 20% minor, leaving only a tiny fraction to the other chord classes. The Jazz standards in the Real Book are more balanced in that respect and feature only 54% of major chords leaving 33% to the minor chords and 6% to the diminished chords.

3.2 The Most Commonly Used Chord Sequences

As argued in section 2.3 we will use length 4 sequences as a starting point for the considerations.

3.2.1 Beatles

In the Beatles songs the prevailing idiom is major chords alternating in fifths (or fourths), appearing in 41% of the songs (see table 2). Typical instances of that would be I-IV-I-IV, but also V-I-V-I. Maybe one could call this the "blues/rock back and forth" idiom. Interestingly, the six highest ranking sequences contain only fifth (+7 semitones) or fourth (+ 5 semitones) chord changes of major chords.

The first chord sequence that features a minor chord ranks at 15. The latter is the second most frequent sequence of the Real Book collection, typically ii-V-I-IV. Moreover, that same chord sequence is also the highest ranking one that does *not* repeat a chord. All the others contain chord changes that add up to 12 or 24 at some point.

3.2.2 Real Book

As was expected, the most common chord sequences in the Real Book follow the circle of fifths idiom (ascending in fourths/descending in fifths). In fact, combinations of that kind of sequence seem to dominate all other sequences. A huge difference to the Beatles can be spotted even at rank 1 of table 3. This chord sequence requires one chord to have a non-diatonic note. So not only do the most widely used Real Book sequences not repeat a chord, they even force some kind of modulation.

3.3 Searching for Known Chord Sequences

Searching the chord sequences for known idioms can help to find evidence for the songwriters' influences. For instance, the opening sequence of the Beatles' *Yesterday*:

$$\text{maj} \stackrel{+11}{\to} \text{min} \stackrel{+5}{\to} \text{maj} \stackrel{+5}{\to} \text{min}$$

	shift							p_i
1	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+7}{\to}$	maj	0.41
2	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+5}{\to}$	maj	0.40
3	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+7}{\to}$	maj	0.31
4	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+7}{\to}$	maj	0.24
5	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+5}{\to}$	maj	0.23
6	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+5}{\to}$	maj	0.23
7	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+2}{\to}$	maj	$\stackrel{+5}{\to}$	maj	0.22
8	maj	$\stackrel{+2}{\to}$	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+5}{\to}$	maj	0.21
9	maj	$\stackrel{+7}{\to}$	maj	$\stackrel{+10}{\to}$	maj	$\stackrel{+7}{\to}$	maj	0.19
				...				
15	min	$\stackrel{+5}{\to}$	maj	$\stackrel{+5}{\to}$	maj	$\stackrel{+5}{\to}$	maj	0.10
				...				
–	maj	$\stackrel{+11}{\to}$	min	$\stackrel{+5}{\to}$	maj	$\stackrel{+5}{\to}$	min	0
				...				

Table 2. Beatles: ranking of length 4 chord sequences with empirical probability p_i of appearing in a song from the Beatles Chord Database

is unique within the Beatles Database, but occurs in 17 songs in the Real Book (rank 29) including the Charlie Parker standard *Confirmation* and *Like Someone in Love* by Johnny Burke.

However, when classifying the chords according to the *one class* set, i.e.

$$\bullet \stackrel{+11}{\to} \bullet \stackrel{+5}{\to} \bullet \stackrel{+5}{\to} \bullet$$

we can find more instances (with differing chord types but identical root differences) in 12 other songs from the Beatles Database, including the early *Do You Want to Know a Secret* and *The Long and Winding Road* on their final album.

4 FUTURE WORK

In order to achieve more generality and to be able to compare different harmony styles we will extend this work to larger and more diverse manually transcribed collections. Where no chord transcription is available we are planning to take advantage of the increasingly reliable methods of chord inference from MIDI data as mentioned above.

Just as for melodies, the meaningful processing of chord sequences depends on their rhythmic structure. We have retrieved relative chord lengths from the Beatles Database's onset and offset times as well as from the beat-aligned notation in the Real Book, which allows us to exploit harmonic rhythm in the future.

The probabilities p_i and q_i complement each other and suggest a two-step generative model for chord sequences in the sense of Temperley's [11] generative model for melodies. Our estimates can be used as parameters in such models, and the information on harmonic rhythm could serve as a prior for advanced duration modelling.

	shift							p_i
1	maj	$\overset{+5}{\to}$	min	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	0.28
2	min	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	0.25
3	maj	$\overset{+0}{\to}$	min	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	0.23
4	min	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	min	$\overset{+5}{\to}$	maj	0.20
			...					
9	min	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	$\overset{+9}{\to}$	min	0.16
			...					
14	maj	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	maj	0.11
15	maj	$\overset{+9}{\to}$	min	$\overset{+5}{\to}$	min	$\overset{+5}{\to}$	maj	0.11
16	dim	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	min	$\overset{+5}{\to}$	maj	0.11
			...					
19	min	$\overset{+5}{\to}$	maj	$\overset{+7}{\to}$	min	$\overset{+5}{\to}$	maj	0.10
			...					
29	maj	$\overset{+11}{\to}$	min	$\overset{+5}{\to}$	maj	$\overset{+5}{\to}$	min	0.08
			...					

Table 3. Real Book: ranking of length 4 chord sequences with empirical probability p_i of appearing in a song from the Real Book

Another source of improvement will be the annotation of section boundaries, marking the singular chord sequences at the beginning and – most importantly – the ending of sections. Besides providing an alignment for chord idioms, this step could also aid other segmentation algorithms.

However, the problem of finding special chord idioms remains, as pure chord sequence frequencies do not explain the phenomenon very well. In linguistics, word frequencies have been used to find collocations, i.e. word pairs that appear in conjuction more often than could be expected from their respective frequencies [4]. We expect to find more meaningful chord sequences by using p_i and q_i in conjunction with those methods. Also, functional chord distances can be calculated from chord sequences in the style of Pachet's substitution rules [7].

Findings in the symbolic domain instantly pose the question if it possible to design equivalent algorithms in the audio domain. The prerequisite to idiom discovery and analysis in this domain is the classification of individual chords or chord classes. We are developing a technique using support vector machines trained with synthesised audio to classify segmented audio by chord type. In preliminary tests, correctness of labels is significantly better than random assignment, with frame-by-frame accuracy of approximately 40%. There is great potential in this method of chord labelling, though in its current state there is room for much improvement. Once accurate chord identification is achieved it is possible to examine idioms in digital audio in much the same way as this paper shows in the symbolic domain.

5 CONCLUSIONS

In this paper we have presented our first approach to harmonic analysis from manually labelled chord symbols aiming at the discovery of chord idioms. We have categorised chords into musically motivated classes with different levels of generalisation, along with a key-independent definition of chord sequences. Chord sequences were retrieved from two collections comprising over 400 songs and ordered according to a proposed frequency measure. Examples show how the method can help to understand the use of chords in the collections and illustrate the musicological relevance of chord idioms. In the future we will explore idioms in both the symbolic and audio domains.

6 ACKNOWLEDGEMENTS

This work is part of the OMRAS2 project funded by EPSRC grant EP/E017614/1.

7 REFERENCES

[1] Merriam-Webster Online Dictionary, April 2007.

[2] J. P. Bello and J. Pickens. A Robust Mid-level Representation for Harmonic Content in Music Signals. In *Proc. ISMIR 2005, London, UK*, 2005.

[3] D. de la Motte. *Harmonielehre*. dtv/Bärenreiter, 1995.

[4] T. Dunning. Accurate Methods for the Statistics of Surprise and Coincidence. *Computational Linguistics*, 19(1):61–74, 1994.

[5] C. Harte and M. Sandler. Automatic Chord Identifcation using a Quantised Chromagram. In *Proceedings of 118th Convention*. Audio Engineering Society, 2005.

[6] C. Harte, M. Sandler, S. A. Abdallah, and E. Gomez. Symbolic representation of musical chords: A proposed syntax for text annotations . In *Proc. ISMIR 2005, London, UK*, 2005.

[7] F. Pachet. Surprising Harmonies. In *International Journal of Computing Anticipatory Systems*, 1999.

[8] J.-F. Paiement, D. Eck, and S. Bengio. A Probabilistic Model for Chord Progressions. In *Proc. ISMIR 2005, London, UK*, 2005.

[9] C. Rhodes, D. Lewis, and D. Müllensiefen. Bayesian Model Selection for Harmonic Labelling. May 2007. MCM Berlin.

[10] D. Sleator and D. Temperley. The Melisma Music Analyzer. web page: www.link.cs.cmu.edu, 2007.

[11] D. Temperley. *Music and Probability*. The MIT Press, January 2007.

[12] various. *The Real Book*. Hal Leonard Corporation, 6th edition, September 2004.