

MASTER THESIS



AUDIO-SYMBOLIC ALIGNMENT OF POPULAR MUSIC
WITH APPLICATION TO AUTOMATIC CHORD ESTIMATION



AUTHOR: DAPHNE ODEKERKEN

SUPERVISORS: ANJA VOLK, HENDRIK VINCENT KOOPS

SECOND EXAMINER: FRANS WIERING

ICA-3827887



Utrecht University

N C C# D Eb E F F# G G# A Bb B Cm C#m Dm Ebm Em Fm F#m Gm G#m Am Bbm Bm

Abstract

Automatic Chord Estimation (ACE) is a fundamental task in Music Information Retrieval (MIR), which has applications in both music performance and in research on other MIR tasks. The ACE task consists of segmenting a music recording or score and assigning a chord label to each segment. ACE has been a task in the MIREX competition for 10 years, but is not yet a solved problem: current methods seem to have reached a glass ceiling. Moreover, many recent methods are trained on a limited data set and consequently suffer from overfitting. In this thesis, I propose DECIBEL¹ (DEtection of Chords Improved By Exploiting Linking symbolic formats), a novel system that utilizes multiple symbolic music representations in addition to audio in order to improve ACE on popular music.

The input for DECIBEL not only consists of the audio file, but also contains a set of MIDI and tab files that are obtained through web scraping and manually matched to the audio file. Given the audio file and matched MIDI and tab files, the system first estimates chord sequences from each file, using a representation-dependent method. For audio files, DECIBEL uses existing state-of-the-art audio ACE techniques: in my experiments, I use the output of the six ACE submissions from MIREX 2017, as well as a commercial state-of-the-art method. MIDI files are first aligned to the audio file, using a Dynamic Time Warping-based method; subsequently, chord sequences are estimated from the re-aligned MIDI files using an algorithm based on template matching. Tab files are first parsed, resulting in untimed chord sequences and then aligned to the audio, using an existing algorithm based on a Hidden Markov Model. In a final step, DECIBEL uses a data fusion method that integrates all estimated chord sequences into one final output sequence.

The main contributions of this thesis are twofold. First, by aligning different symbolic formats to audio, DECIBEL automatically creates a heterogeneous harmonic representation that enables large-scale cross-version analysis of popular music. Second, my results show that DECIBEL’s data fusion method significantly improves each of the seven evaluated state-of-the-art audio ACE methods in terms of estimation accuracy. Exploiting the musical knowledge that is implicitly incorporated in MIDI and tab files, it breaks the observed glass ceiling, without requiring a lot of additional training, thereby prohibiting further overfitting to the existing chord annotations.

¹An implementation of DECIBEL is available on <https://github.com/Daphne0/DECIBEL>.

Contents

1	Introduction	1
1.1	Stagnation and subjectivity	2
1.2	Research goal	3
1.3	Related work	3
1.4	Thesis outline	4
2	Musical background	5
2.1	An introduction to pitch	5
2.1.1	Pitch in physics: waves and frequencies	5
2.1.2	Twelve-tone equal-temperament	6
2.1.3	Note names and octave numbers	7
2.2	Music notation	8
2.3	Intervals	9
2.4	Chords	10
2.4.1	Chord notation and representations	13
2.5	Music representations	15
2.5.1	Audio representation	15
2.5.2	MIDI representation	19
2.5.3	Tab representation	21

2.6	Summary of musical background	23
3	Framework of proposed system “DECIBEL”	25
3.1	DECIBEL’s framework in a nutshell	25
3.2	Collection of data set	27
3.3	Performance evaluation	30
4	Automatic Chord Estimation on audio	33
4.1	Related work on audio ACE	33
4.1.1	Feature extraction for audio ACE	35
4.1.2	Audio ACE models	36
4.2	Selected systems	37
4.3	Evaluation of audio ACE	38
4.4	Conclusion	41
5	Automatic Chord Estimation on MIDI	43
5.1	MIDI-to-audio alignment	44
5.1.1	Related work on MIDI-to-audio alignment	44
5.1.2	Selected system	49
5.1.3	Evaluation of MIDI-to-audio alignment	52
5.2	Chord estimation on MIDI	52
5.2.1	Related work on chord estimation on MIDI	53
5.2.2	Implementation of CASSETTE	54
5.2.3	Evaluation of CASSETTE	56
5.3	MIDI file selection	57
5.4	Conclusion	59

6	Automatic Chord Estimation on tabs	61
6.1	Tab parsing	62
6.1.1	Tab parsing methodology	62
6.1.2	Tab parsing evaluation	65
6.2	Jump Alignment	66
6.2.1	Preprocessing	66
6.2.2	Hidden Markov Model	66
6.2.3	Jump Alignment	68
6.3	Tab file selection	70
6.4	Conclusion	72
7	Data fusion	73
7.1	A rich harmonic representation	74
7.2	Related Work on data fusion	76
7.3	Data fusion experiments	77
7.4	Data fusion results	79
8	Conclusion	85
8.1	Contributions	85
8.1.1	Rich harmonic representation	86
8.1.2	Improvement of state-of-the-art audio ACE systems	86
8.1.3	Extended Isophonics data set	86
8.1.4	Python implementations	87
8.2	Future work	87
8.3	Acknowledgements	88

A	Data set	89
B	Alignment listening test results	95
C	Tabs in parsing evaluation set	97
D	Tabs in forward and backward training set	99
E	Tukey's HSD test results	101
E.1	CHF	101
E.2	CM2	102
E.3	JLW1	103
E.4	JLW2	104
E.5	KBK1	105
E.6	KBK2	106
E.7	WL1	107

Chapter 1

Introduction

Automatic Chord Estimation (ACE)¹ is a fundamental problem in Music Information Retrieval. The ACE task is concerned with estimating chords in audio recordings or symbolic representations of music. Basically, ACE segments a song in such a way that the segment boundaries represent chord changes and each segment has a chord label. This is typically represented by a sequence of ⟨start time, end time, chord label⟩ 3-tuples.

Figure 1.1 shows an example of such a chord sequence. We see that the first 2.6 seconds of this song do not contain any chords, hence the no-chord symbol N. This no-chord segment is followed by a 8.8 second segment in which an E major chord sounds, followed by an A major chord, etcetera.

```
0.000000 2.612267 N
2.612267 11.459070 E
11.459070 12.921927 A
12.921927 17.443474 E
17.443474 20.410362 B
20.410362 21.908049 E
21.908049 23.370907 E:7/3
23.370907 24.856984 A
24.856984 26.343061 A:min/b3
26.343061 27.840748 E
27.840748 29.350045 B
```

Figure 1.1: Chord sequence annotation for the beginning of the Beatles song *I Saw Her Standing There*. This is a sequence of ⟨start time, end time, chord label⟩ 3-tuples.

A chord can be defined as multiple notes that sound simultaneously. As we will see in Section 2.4.1, any chord can efficiently be represented by a text string (for example the chord labels in Figure 1.1: E, A, B, E:7/3 and A:min/b3). The progression of chords through time defines the harmonic structure of a piece of music. Therefore, the estimation

¹ACE is also referred to as *Chord Recognition*, *Automatic Chord Transcription* or *Automatic Chord Detection*.

of chords in a music piece has many applications in both music performance and in Music Information Retrieval (MIR), as described below.

For music students, it is common to play along with songs on on-line streaming services like YouTube (Stowell and Dixon, 2011). ACE can make this easier for the less experienced musician. Chordify² is a web-service that uses ACE techniques to display the chords of any audio file to the user. From the fact that Chordify is used by 1.5 million users every month (Bountouridis et al., 2016), we can conclude that there is great interest in systems using ACE.

Chord sequences have also been used by the MIR research community in high-level tasks such as cover song identification (Gómez and Herrera, 2006; Ellis and Poliner, 2007; Maršík et al., 2017), key detection (Papadopoulos and Tzanetakis, 2012), genre classification (Ajoodha et al., 2015), lyric interpretation (Kolchinsky et al., 2017) and audio-to-lyric alignment (Mauch et al., 2012). To conclude, a well-performing ACE system, which estimates reliable chord sequences, would be of great use for both music performers and MIR researchers.

1.1 Stagnation and subjectivity

Fujishima (1999) was the first who considered ACE as a problem on its own and since this pioneering publication, many researchers have worked on the task. Despite these efforts in the past two decades, ACE is not yet a solved problem.

ACE has been a task in the annual benchmarking evaluation Music Information Retrieval Evaluation eXchange (MIREX) since 2008. The main evaluation measure is Weighted Chord Symbol Recall (WCSR), which reflects the proportion of correctly labeled chords in a data set, weighted by the total length of the data set. State-of-the-art ACE methods yield WCSRs of around 80%, given a chord vocabulary of major and minor chords³. However, from recent MIREX results, Humphrey and Bello (2015) and Scholz et al. (2016) observe a stagnation in ACE performance. Scholz et al. (2016) give two suggestions to recover from this stagnation: ACE methods should either make better use of musical expert knowledge, or they should use new techniques, for example deep learning, instead of relying on the small set of commonly used techniques.

Besides, Humphrey and Bello (2015) throw light on another issue of ACE: chord annotations are inherently subjective. Even between musically trained human annotators, there can be a discrepancy of over 15% in the chord annotation of a song. This subjectivity matter was earlier identified by Ni et al. (2013), who asked 5 musicians to each annotate the same 20 songs and show a 10% discrepancy between the annotations. Both Humphrey and Bello and Ni et al. dispute the currently common practice to evaluate ACE by comparing the results to a single reference annotation. Ni et al. (2013) even claim that modern ACE systems have started to overfit the MIREX data set, mimicking

²<https://chordify.net/>

³http://www.music-ir.org/mirex/wiki/2017:Audio_Chord_Estimation

the subjective aspects of MIREX’s reference annotations. Although neither paper presents a practical, scalable solution to the subjectivity issue, it becomes clear that subjectivity in ACE is an important issue that is often overlooked in existing research.

In conclusion, there is need for a new strategy that overcomes existing stagnation in ACE without further overfitting to existing (subjective) data sets. In this thesis, I propose a novel method that uses MIDI and tab files, which implicitly contain crowd-sourced musical knowledge, to significantly improve existing state-of-the-art methods, requiring a minimal amount of additional training.

1.2 Research goal

The aim of this research is to show that audio ACE can be improved by exploiting symbolic representations of popular music. For this purpose, I designed and implemented DECIBEL (DEtection of Chords Improved By Exploiting Linking symbolic formats). DECIBEL is a novel system that exploits multiple heterogeneous symbolic music representations for improving ACE. MIDI and tab files can be considered as crowd-sourced note and chord transcriptions respectively. By using these symbolic representations, DECIBEL implicitly integrates musical knowledge into existing ACE methods, which may be a strategy to overcome the stagnation issue identified in the previous section. As DECIBEL only relies to a small extent on training on reference annotations, my method prohibits further overfitting to existing data sets. To evaluate DECIBEL, I compare its performance to state-of-the-art ACE systems submitted in the MIREX competition of 2017 and show that DECIBEL’s results are a significant improvement.

1.3 Related work

Integrating symbolic and audio formats in a chord estimation setting is not a completely new idea. Both MIDI and tab files were used for detecting chords in earlier work, as we will see in this section.

Ewert et al. (2012) introduce a cross-version analysis framework for comparing harmonic analysis results from different musical domains. After collecting a MIDI file for each audio file in a 112-song subset of the Isophonics (Mauch et al., 2009; Harte, 2010) data set, they use two state-of-the-art chord recognition methods for MIDI data and align each MIDI file to the corresponding audio recording. This way, they create a harmonic representation for each of the 112 songs, which contains three chord label sequences: the ground truth labels and the re-aligned outputs that were obtained by the two MIDI chord recognition systems. They show that this harmonic representation can be used for quantitative evaluation of MIDI chord recognition methods, using annotations for corresponding audio recordings. In addition, by visualizing this harmonic representation, they demonstrate how it can be used for qualitative error analysis of automatically generated chord labels, and by that contributes to the understanding of an ACE algorithm’s behavior and the properties of

the underlying music material. The research by Ewert et al. lays the foundation for the work proposed in this thesis, in which I expand the harmonic representation with chord labels from multiple MIDI and tab files for each audio recording and show how this enriched harmonic representation can be used to improve ACE.

The integration of tab files and audio with respect to chord estimation was earlier researched by McVicar and De Bie (2010) and McVicar et al. (2011a,b). In these three papers, they show that a HMM-based system for audio ACE can be significantly improved by incorporation of external information from guitar tabs. In a preprocessing step, they obtain a set of tab files by a web scrape and consequently parse them. In the resulting format, to which they refer as Untimed Chord Sequences (UCSS), only the chord labels and line information of each tab is retained. As a next step, they align each UCS to the corresponding audio file. For this purpose, the authors introduce four variations on the traditional Viterbi algorithm. The most promising variation is Jump Alignment, which aligns the UCS to the audio file, thereby allowing jumps from the end of any annotation line to the beginning of any line. I implemented Jump Alignment as part of the DECIBEL system and will explain the implementation more thoroughly in Section 6.2.

1.4 Thesis outline

The remainder of this thesis is structured as follows: Chapter 2 provides the required music theory and terminology information in order to understand the concepts that will be used in the remainder of this study. In Chapter 3, I describe DECIBEL's framework. As we will see in this chapter, DECIBEL consists of four subsystems: three of these subsystems compute chord labels in their own representation-specific way and the fourth subsystem combines these results in a data fusion step. The audio, MIDI and tab subsystems are described in Chapters 4, 5 and 6 respectively. DECIBEL's data fusion strategy is described in Chapter 7. Finally, I will present my conclusions in Chapter 8.

Chapter 2

Musical background

This chapter gives an introduction to pitch, music notation, intervals and chords. The material in Section 2.1 to 2.4 is based on Taylor (1989) and (Müller, 2015, Chapter 1 and 5). Section 2.5 describes the three music representations we consider in this research: audio, MIDI and tabs.

2.1 An introduction to pitch

Music consists of tones, and each tone has some properties, for example its duration, start time and **pitch**. Pitch refers to the degree of highness of sound. If you hit a key on the left side of the piano and subsequently strike a key on the right, the second produced tone will be higher pitched than the first one. Similarly, men's voices are generally lower pitched than women's voices and a tuba sounds lower than a piccolo.

2.1.1 Pitch in physics: waves and frequencies

From a physical point of view, sound is generated by vibrating objects, for example the string and soundboard of a violin or the vocal cords of a singer. These vibrations cause displacements and oscillations of air molecules, resulting in local regions of compression and rarefaction. This alternating pressure travels through the air as a **longitudinal wave**, from its source to a perceiver.

The change in air pressure at a certain location can be graphically represented by a **waveform** of the sound. A waveform plots the deviation of the air pressure from the average air pressure over time. This is illustrated in Figure 2.1.

If the points of high and low air pressure repeat in an alternating and regular fashion, the resulting waveform is called periodic. The **period** is defined as the amount of time that is required for completing a cycle. The **frequency** is the reciprocal of period, and




Figure 2.1: Illustration of sound waves and the waveform, from (Müller, 2015, Chapter 1). A waveform plots the deviation from the average air pressure over time.

is measured in Hertz (Hz). For example, a period of 2.5 ms corresponds with a frequency of $\frac{1}{0.0025 \text{ s}} = 400 \text{ Hz}$. The higher the frequency of a sinusoidal wave, the higher the corresponding tone sounds.

However, real-world sounds, such as the tones that are produced by a musical instrument, are much more complex than a pure tone resulting from a single sinusoid. A musical tone can be described as a superposition of sinusoids, each with their own frequency. A **partial** is any of the sinusoids by which a musical tone is described. The frequency of the lowest partial present is called the **fundamental frequency** of the sound. The pitch of a musical tone is usually determined by the fundamental frequency.

Two tones with fundamental frequencies in a ratio equal to any power of two, are perceived as similar. All tones with this kind of relation can be grouped in the same **pitch class**. The distance between one musical tone and another tone with half or double its fundamental frequency is called an **octave**. For example, a tone with a frequency of 440 Hz sounds similar to a tone with a frequency of 220 Hz. These tones are an octave apart and belong to the same pitch class.

2.1.2 Twelve-tone equal-temperament

We have seen that pitch and frequency are closely related, but musicians typically do not specify the height of a tone in Hertz. Instead, the space of all different pitches is discretized using a **tuning system**. Although many different tuning systems have been suggested and used in history, the standard tuning in modern Western music, which is the tuning used in this research, is **twelve-tone equal-tempered tuning**. This tuning system is the standard system used as basis for tuning the piano.

In the twelve-tone equal-tempered tuning system, an octave is divided into twelve scale steps - often simply referred to as tones. The fundamental frequencies of these steps are equally spaced on a logarithmic frequency axis, as the human perception of pitch is logarithmic in nature. This means that the frequency ratio of two subsequent scale steps is constant and equals $2^{\frac{1}{12}} \approx 1.059463$. The distance, or **interval**, between two subsequent scale steps is called a **semitone**. In other words, if we multiply the frequency of an arbitrary pitch by $2^{\frac{1}{12}}$, this pitch is raised by a semitone.

These semitones can be further divided into **cents**. By definition, each octave is divided into 1200 cents, so a semitone consists of 100 cents. The frequency ratio of two subsequent cents equals $2^{\frac{1}{1200}} \approx 1.0005777895$. Note that both the cent and the semitone are logarithmic measures of distance between pitches.

2.1.3 Note names and octave numbers

In Section 2.1.1, we have seen that two tones that differ in pitch by one or more octaves can be grouped in the same pitch class and in Section 2.1.2 we learned that the twelve-tone equal-tempered tuning system divides the octave into twelve steps of one semitone. We can therefore deduce that there are twelve pitch classes.

In modern Western music notation, each pitch class has a **pitch class name** or note name, which consists of a letter and possibly an accidental. There are seven pitch classes which can be named by just one letter from $\{A, B, C, D, E, F, G\}$. These pitch classes correspond to the white keys of the piano, as illustrated in Figure 2.2.




Figure 2.2: An octave on the piano. Pitch classes corresponding to the white keys of the piano can be named by a single letter. The remaining five pitch classes are named by a combination of a letter and an accidental.

The remaining five pitch classes are named by a combination of a letter and an **accidental**. An accidental raises or lowers the corresponding note. The most common accidentals are the sharp (\sharp) and flat (\flat), which respectively raise and lower the note with a semitone. So the black key on the piano positioned between the C- and D-keys produces a note that can be denoted by either $C\sharp$ or $D\flat$. Similarly, the other “black-key” pitch classes can be referred to by either $\{D\sharp, F\sharp, G\sharp \text{ and } A\sharp\}$ or $\{E\flat, G\flat, A\flat \text{ and } B\flat\}$. We see that two different names can refer to the same pitch class. This phenomenon is called **enharmonic equivalence**.

Following **Scientific Pitch Notation**, a pitch is not only specified by the pitch class name, but also by an **octave number**. The higher the octave number, the higher the pitch. For instance, an A1 sounds an octave higher than an A0, which is the lowest tone that can be produced by most pianos. The note A4 has a frequency of 440 Hz in modern twelve-tone equal-tempered tuning and is used as a reference note for tuning.




Figure 2.3: The staff, with its lines and spaces




Figure 2.4: Notes on the staff, with their names described above. Here we see that when using different clefs, the same note is notated at a different (vertical) position on the staff.

Note that the distance between notes denoted by two subsequent letters is not always the same. For example, D4 and E4 differ by two semitones: in Figure 2.2 we see that there is a distance of two piano keys, so two semitones, between D and E. On the other hand, E4 and F4 are just one semitone apart: there is no black key on the piano between these notes.

2.2 Music notation

Music is notated on a **staff**, which consists of five parallel horizontal lines, counted from the bottom, see Figure 2.3. Notes can be notated on the lines or between them. The pitch of a note is determined by its vertical position on the staff, combined with the **clef**. The two most common clefs are the treble clef or ‘G’ clef (G) and the bass clef or ‘F’ clef (F). The working of these clefs is illustrated in Figure 2.4: if there is a treble clef at the beginning of the staff, this means that the note on the second line is G4. By way of illustration, this note is colored red. When using the bass clef, the note that is placed on the fourth line is F3 - colored blue.

If notes lie above or below the limits of the staff, short additional lines called **ledger lines** are used. We see for example that we need one ledger line for the notation of C4 in the treble clef as well as in the bass clef. As ledger lines deteriorate the readability of sheet music, it is common to use the treble clef for high notes (C4 and higher) and use the bass clef for lower notes.



Figure 2.5: The eight intervals of the major scale. Note that the distance between the third (E) and fourth (F) note and between the seventh (B) and eighth (C) note is only one semitone.

2.3 Intervals

Now that we have become acquainted with tones, pitches and their notation, it is time to introduce the notion of intervals. An **interval** is the distance between two pitches. We already know that, assuming twelve-tone equal-tempered tuning, we can divide the octave in twelve semitones. Based on the notion of semitone, we can specify other intervals that are used in Western music theory.

Interval names consist of a **number** and a **quality**. In counting the number, both notes are included. For example when determining the interval from C₄ to F₄, we count four notes (C₄, D₄, E₄, and F₄), so this interval is called a fourth. Similarly, the interval from E₄ to C₅ is a sixth. The interval from one note to another note with exactly the same pitch is a **unison**.

Traditionally, intervals are named on the basis of the **major scale**, which consists of seven notes and an eighth note one octave apart from the first note. In the major scale, there is only one semitone difference between note 3 and 4 and between note 7 and 8, while there are two semitones difference between all other subsequent notes. An example of a major scale is C major: C - D - E - F - G - A - B - C. In Figure 2.5 we see the full names of all intervals between C, which is the root note of the scale, and each of the other notes. In each major scale, the intervals unison, fourth, fifth and octave get the quality “perfect”, while the intervals second, third, sixth and seventh get the quality “major”.

Using only perfect and major intervals, we cannot express all different intervals. For instance, we have no name yet for an interval of three semitones. We can solve this by raising or lowering the upper note of the interval with a semitone, following these rules:

- Given a perfect or major interval: if either the upper note is raised a semitone or the lower note is lowered a semitone, the interval becomes **augmented**;
- Given a major interval: if either the upper note is lowered a semitone or the lower note is raised a semitone, the interval becomes **minor**;
- Given a perfect or minor interval: if either the upper note is lowered a semitone or the lower note is raised a semitone, the interval becomes **diminished**.

Figure 2.6 shows all possible major, minor, augmented and diminished intervals in (or just beyond) the octave, with their distance in semitones. Note that two different interval names can refer to the same difference in semitones, e.g. both the augmented fourth and the diminished fifth consist of six semitones. These intervals are enharmonically equivalent.

2.4 Chords

A **chord** can be loosely defined as a group of tones sounding at the same time. Most researchers agree that a chord should consist of tones from at least three distinct pitch classes.

Chords that consist of tones from three pitch classes are called **triads**. In Western music, most triads can be stacked in thirds and consist of the note on which the triad is based (**root**), plus the third and the fifth above it. The root note determines the name, while the quality of the other intervals determines the type of chord. For example: a C with a major third and a perfect fifth forms a C major chord, and a D with a minor third and a perfect fifth is a D minor chord. It is possible to extend chords by stacking more thirds upon them. **Seventh** chords consist of a “normal” triad and an added seventh, forming a tetrad. The most common seventh chord is the dominant seventh, made out of a major tetrad and a minor seventh. Figure 2.7 summarizes the most used triad and tetrad chords in Western music.

For one chord, there are many variations possible. First, there exist multiple **inversions** for each interval. When the chord’s lowest note is its root (like in the examples of Figure 2.7), the chord is said to be in **root position**. When the lowest note is the third, for example the E in a C major chord, this chord is in **first inversion**. When the lowest note is the fifth, the chord is in **second inversion**. A seventh chord can even be in **third inversion** if the seventh is the lowest note. Second, notes of the same pitch class may be doubled, for example in a C major chord consisting of C3, E3, F3 and C5. This is called **octave doubling**. Third, the notes of a chord may not be played simultaneously, but after each other. In this case, we speak of a **broken chord**.

The concept of harmony and chords is enriched by the existence of **non-harmonic tones**. These tones are not part of the chord. Instead, they are used to create a smooth melody line, to prepare the transition to another chord or to add a dissonant element, creating musical tension. Though non-harmonic tones certainly contribute to the beauty of music, they provide a challenge for ACE systems as it is difficult to automatically determine whether a given note is either harmonic or non-harmonic.




Figure 2.6: Major, minor, augmented and diminished intervals, with their distance in semitones

The figure displays eleven rows of musical notation, each representing a different chord type. Each row consists of five measures. The first measure shows the full chord (triad or tetrad) with its key signature. The subsequent measures show individual notes with labels above them: 'Root note', 'Major third', 'Minor third', 'Perfect fifth', 'Diminished fifth', or 'Augmented fifth'. The chords and their constituent notes are as follows:

- Major:** Root note, Major third, Perfect fifth
- Minor:** Root note, Minor third, Perfect fifth
- Diminished:** Root note, Minor third, Diminished fifth
- Augmented:** Root note, Major third, Augmented fifth
- Major seventh:** Root note, Major third, Perfect fifth, Major seventh
- Minor seventh:** Root note, Minor third, Perfect fifth, Minor seventh
- Dominant seventh:** Root note, Major third, Perfect fifth, Minor seventh
- Diminished seventh:** Root note, Minor third, Diminished fifth, Diminished seventh
- Half-diminished seventh:** Root note, Minor third, Diminished fifth, Minor seventh
- Minor major seventh:** Root note, Minor third, Perfect fifth, Major seventh
- Augmented major seventh:** Root note, Major third, Augmented fifth, Major seventh

Figure 2.7: The most common chords based on triads and tetrads

2.4.1 Chord notation and representations

There exist various notations for chords, which can differ between and even within genres. (Taylor, 1989) In baroque music for example, the chord notation consists of a bass line with figures written underneath the notes. This notation is called **figured bass** and is illustrated in Figure 2.8. The figures represent the intervals that should be played above the bass note. For example: the first bar is the figured bass notation for a C major chord in root notation, and can be played like the second bar in Figure 2.8. Similarly, the figured-bass notation in the third bar means that we should add the third and the sixth to the bass note. The bass note is an E, so we add a G and C to the bass note, resulting in a C major chord in first inversion - as written out in the fourth bar. The fifth and sixth bar represent a second-inversion C major chord in figured bass and full notation respectively.




Figure 2.8: Three inversions of the C chord in figured bass notation

In Classical Harmony Analysis, chords are studied and notated in relation to the current key. This notation is called **Roman numeral notation** and is, unsurprisingly, characterized by Roman numerals under the chords. These Roman numerals indicate the scale degree on which the chord is built, as illustrated in Figure 2.9. In this figure, we see seven chords in the key of C major. The first chord is a C major chord. The interval between the key C and the root note of the chord (C) is a unison, so we need the Roman equivalent of the digit “1”. As a rule, major chords are capitalized, so the correct Roman numeral for this chord in this key is “I”. The second chord is a D minor chord. The key (C) and the root note of the chord (D) are a second apart and the chord is in minor, so the corresponding numeral is “ii”. The final bar in our example piece is a diminished B chord. Diminished chords are indicated with a small “0” after the Roman numeral. Similarly, augmented chords (not present in this example) are indicated with a small “+” after the Roman numeral. Note that the Roman numeral that a chord gets, is dependent on the key in which it occurs. In Figure 2.10, we see for example that a G major chord gets the Roman numeral “V” in the key of C major, while “I” is the right numeral in the key of G major.




Figure 2.9: Seven chords in the scale of C major in Roman numeral style notation notation

Yet another chord representation is used in **jazz and popular music**. These genres are characterized by improvisation in performance. The basic harmonies are generally quite simple, so the main purpose of the chord representation is to be easy to read and interpret. As we will see in Section 2.5.3, this chord representation is common in tabs (an




Figure 2.10: Roman numerals are dependent on the current key

alternative for regular sheet music). Figure 2.11 shows some variations of the C chord and their chord notation in jazz and other forms of popular music. The letter name of the root of the chord is shown as a capital letter. A chord is assumed to be major unless stated otherwise. Minor chords get a “m” after the chord letter; “+” or “aug.” stands for an augmented chord and a chord with a “o” or “dim.” is diminished. It is unusual to specify the bass note, so a chord without bass indication can be in any inversion. In the rare cases where the bass note is specified, it is notated with a slash after the letter name of the root of the chord, like in the last two bars of our example. It is important to notice that there exist multiple variations of chord notation in jazz and pop music, and some of them cause ambiguity issues. For example: some musicians use the symbol Δ for major chords, so $C\Delta$ is a C major chord and $C\Delta^7$ is a C major seventh chord. Others use the Δ symbol as a synonym for major seventh chords, so they would notate the C major seventh chord as $C\Delta$. To interpret chords like $C\Delta$ correctly, one should therefore be well aware which variation of notation is applied.



Figure 2.11: Major, minor, augmented and diminished C chords in root position; Major C chords in first and second inversion

Each of the three chord notations is suitable for the genre or field of study in which they are used. However, none of them is well-suited for the chord annotations which we need in order to train and test ACE systems: they are either hard to write in flat text, key-dependent or ambiguous. **ACE chord annotations** require a unambiguous, context-independent notation that is easy to write and intuitive to interpret. Harte et al. (2005) propose a chord grammar in which millions of chords can be defined unambiguously, whereas they give a succinct short-hand notation for the most common chords. Chords names are based on their root note, independent of the present key. In this notation, the C minor chord can either be represented by its components, like $C: (b3, 5)$, or by a shorthand string: $C: \text{min}$. We can also compose more complex chords using the component notation: $A: (3, 5, b7, 9)$ is a dominant ninth chord, consisting of A, $C\#$, E, G and B. The asterisk is used as “omit symbol”, so $D: \text{maj}7(*3)$ would be equivalent to $D: (5, 7)$. The bass can be specified with a slash, followed by the interval from root to bass note. A D diminished seventh chord in third inversion would be represented as $D: (b3, b5, bb7)/bb7$ or $D: \text{dim}7/bb7$. This chord grammar has become the standard notation for ACE reference annotations, including the Isophonics annotations (Mauch et al., 2009; Harte, 2010) which I use in this project.

2.5 Music representations

Apart from sheet music, there exist multiple other music representations. Each of them has its own way of storing pitch or chord information. In this section, we investigate the three music representations that are used in DECIBEL: audio, MIDI and tabs.

2.5.1 Audio representation

Sound is generated by vibrations and travels through the air as a longitudinal wave, which can be graphically represented by a waveform (see Section 2.1.1). This waveform plots the deviation of the air pressure from the average air pressure over time. A **digital** audio representation (e.g. a .wav or .mp3 file) is an approximation of this waveform, in which the sound wave of the audio signal is digitized by sampling and quantization. (Müller, 2015, Chapter 2).

Sampling refers to the process of reducing a continuous-time (CT) signal to a discrete-time (DT) signal, which is defined only on a discrete subset of the time axis. The typical sampling rate for CD recordings is 44.1 kHz, so a CD has 44100 samples per second.

In **quantization**, the continuous range of possible amplitudes is replaced by a discrete range of possible values. For CD recordings, a 16-bit coding scheme is used, which allows representation of $2^{16} = 65536$ possible values.

Although the tones that are present at a given time cannot be read directly from the waveform, there exist methods to estimate them. Remember from Section 2.1.1 that a musical tone is basically a superposition of sinusoids, each with their own frequency. Using **Fourier analysis**, we can decompose our signal (the digitized waveform) into the sinusoids it consists of - and their frequencies. In the remainder of this section, I give a short summary of the discrete Fourier transform, short-time Fourier transform and Constant-Q transform. A detailed explanation of Fourier analysis is beyond the scope of this thesis. For more information about Fourier transforms, I refer the reader to Müller (2015, Chapter 2).

The **Fourier transform** converts a signal that depends on time into a representation that depends on frequency. The formula for the discrete Fourier transform is given in Equation 2.1. X_k is the k th Fourier coefficient, which is the amount of frequency k that is present in a signal x . The signal x consists of N time samples.

$$X_k = \sum_{n=0}^{N-1} x(n) \exp\{-2\pi i k n / N\} \quad (2.1)$$

The magnitude of the Fourier transform tells us about the signal's frequency content: if X_k is high for a frequency k , then this frequency k is important in the signal. However, we cannot infer at which time the frequency content occurs. This is illustrated in Figure 2.12: each of the three signals at the top consist of a low note (3Hz) that sounds for 5 seconds

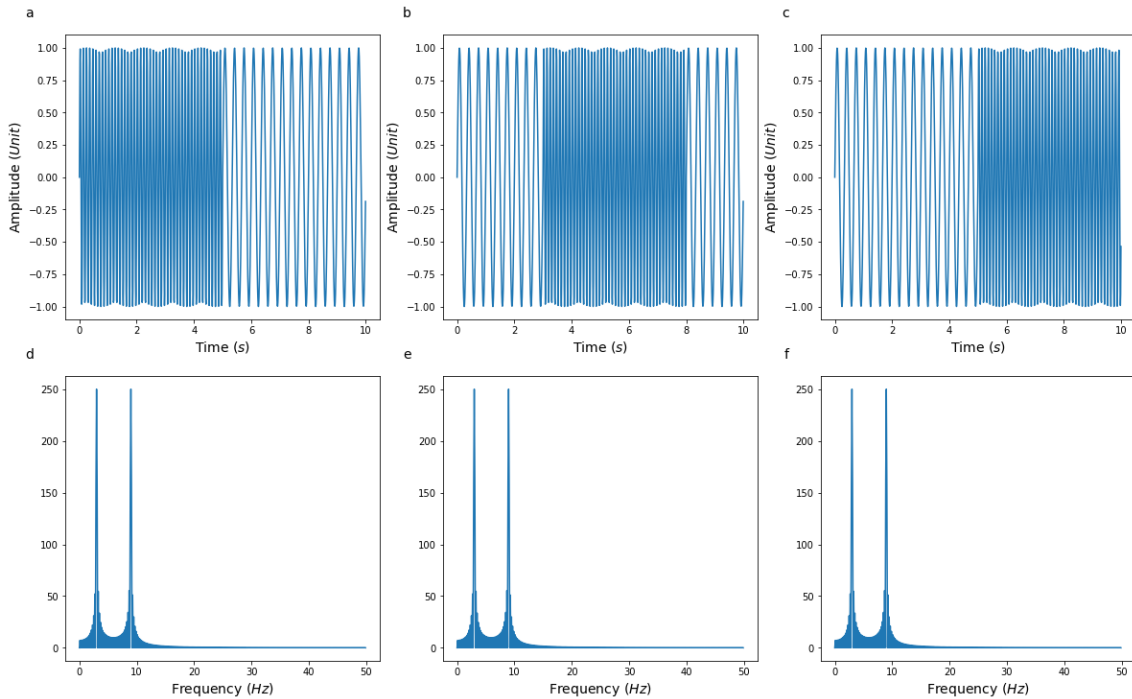


Figure 2.12: Three different signals (a, b, c) with the same frequency distribution yield identical Fourier transforms (d, e, f)

and a higher note (9Hz) that sounds for the other 5 seconds. However, signal *a* starts with the higher note, while in signal *b* this note sounds in the middle and in signal *c* the high note is placed at the end. Though the signals would sound different, the magnitude of the Fourier transform is exactly the same: we see peaks at 3 and 9 Hz.

The **short-time Fourier transform** (STFT) (Gabor, 1946) offers a solution for this problem. The STFT considers only a small part of the signal, which is dependent on the window function w . w is a function with $w(n) \in \mathbb{R}$ if $n \in [0, N - 1]$ and $w(n) = 0$ otherwise. $N - 1$ is the **window size**. The window is shifted every H samples. $H \in \mathbb{N}$ is the **hop size**. In general, a smaller hop size gives more precise results, but is computationally more expensive than a large hop size. The choice for the perfect hop size value is therefore dependent on the application.

$$X_{m,k} = \sum_{n=0}^{N-1} x(n + mH)w(n) \exp\{-2\pi i kn/N\} \quad (2.2)$$

The formula for the STFT is given in Equation 2.2. In Figure 2.13, we see an example of the STFT with a rectangular window of 200 samples (2 seconds with a sampling rate of 100 Hz). The hop size is 200. These window and hop sizes are way too large for real-life applications, but are chosen in order to be visible in the figure.

We can visualize the intensity of frequencies over time in 2D using a **spectrogram**. A spectrogram is the squared magnitude of the STFT:

$$S(m, k) = |X_{m,k}|^2 \quad (2.3)$$

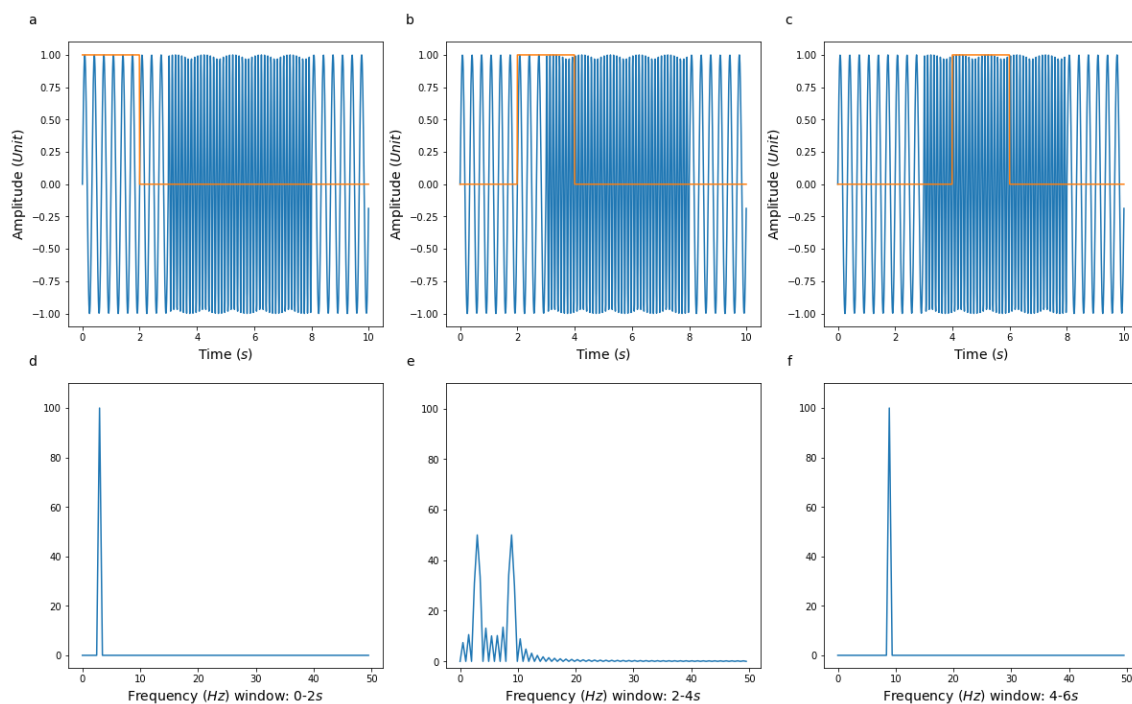


Figure 2.13: Short-time Fourier transform with shifting window: **(a, d)** Window centered at $t = 1s$ **(b, e)** Window centered at $t = 3s$ **(c, f)** Window centered at $t = 5s$

Figure 2.14 shows the spectrogram of *I Saw Her Standing There* by The Beatles. For this spectrogram, we chose a sampling rate of 22050Hz, window size of 2048 samples and hop size of 512 samples. Each point in the spectrogram represents the frequency intensity at a given time point. The warmer the color, the higher the frequency.

The STFT is a common method to extract feature information from a signal. Nevertheless, this method has some problematic properties: first, one needs to specify a window function. The window must be large enough to capture the lowest frequencies. On the other hand: the larger the window, the lower the time resolution. Another disadvantage of the STFT is that the frequencies calculated by the STFT are separated by a constant




Figure 2.14: STFT spectrogram of *I Saw Her Standing There* by The Beatles. Note that the y -axis represents the frequency on a linear scale.




Figure 2.15: CQT spectrogram of *I Saw Her Standing There* by The Beatles. In this spectrogram, the frequencies are placed logarithmically on the y -axis, because the octave numbers increase linearly.

frequency difference, while we have seen that the frequencies of notes of a scale of Western music increase exponentially. Therefore, the frequencies from the STFT do not map directly to the frequencies of music notes.

In the **Constant-Q transform** (Brown, 1991), the window size is not constant, but dependent on the coefficient k : the window N_k grows with higher frequencies. Also, the frequency filters f_k are not spaced linearly (like in the STFT), but logarithmically: the k 'th filter is $f_k = (2^{1/b})^k f_{min}$.

$$X_{m,k} = \frac{\sum_{n=0}^{N_k-1} x(n + mH)w(k, n) \exp\{-2\pi ikn/N_k\}}{N_k} \quad (2.4)$$

The formula of the Constant-Q transform is given in Equation 2.4. Note the differences with Equation 2.2: the window function now has two parameters (k and n); the window size N_k is dependent on k and we normalize by N_k to compensate for high values at high frequencies. When choosing the right values for f_{min} (the lowest frequency we can detect) and b (the number of bins per octave), the Constant-Q transform maps directly to frequencies of musical notes.

In Figure 2.15, we see the spectrogram of the same Beatles song, but now calculated from the Constant-Q transform with $f_{min} = 65.4\text{Hz}$ (the note C2) and $b = 12$. Note that the frequencies are placed logarithmically on the y -axis because the octave numbers increase linearly, so it is easier to see which pitches sounded at each time instance in the song. Almost every pitch has a strictly positive intensity: the lion's share of the spectrogram's colors is not dark blue. Many of those pitches were not played intentionally by The Beatles' band members, but are caused by partials (any of the sinusoids of which a complex tone is composed, see Section 2.1.1).

For most MIR tasks concerning pitches, octave information can be discarded. This is also the case for ACE: the octave numbers are irrelevant in determining the chord from the pitches of which it is composed; we only need the pitch classes. **Chroma features**




Figure 2.16: CQT chromagram of *I Saw Her Standing There* by The Beatles. In this representation, octave information is discarded, but pitch class information is retained.

aggregate all spectral information that relates to a given pitch class into a single coefficient. There exist many variations of Chroma features, and they can be calculated from both the STFT and the Constant-Q transform. Basically, a Chroma feature is a 12-dimensional vector that can be obtained by summing all frequencies. Figure 2.16 shows a visualization of the chroma features of our example song. Here, we see for example that the pitch classes E and G# have high chroma values in the beginning of the song. Indeed, the song starts with a E major chord.

2.5.2 MIDI representation

MIDI is an abbreviation for Musical Instrument Digital Interface. It is a protocol which allows electronic instruments and other digital musical tools to communicate with each other, by sending event messages. Such an event message can for example instruct a synthesizer to start playing a certain note (by a note on event), stop playing a note (note off event) or change to another instrument sound (program change event). A MIDI file is a sequence of MIDI messages, organized in a specific format.

The information stored in MIDI files is therefore fundamentally different from the information stored in audio files: as we have seen in the previous section, audio files represent the waveform of a sound. Conversely, a MIDI file stores a list of instructions for a synthesizer, just like a musical score in a way stores instructions for a musician. The MIDI file itself does not contain any audio signals, but audio can be synthesized based on the MIDI event messages. MIDI files can thus be considered as a compact way to store a musical score and form therefore a symbolic music representation. Compared to audio, it is way easier to extract note information from MIDI. This makes MIDI particularly interesting for research in MIR and musicology.

Another difference between MIDI and audio is the file size: as MIDI stores music on a note level instead of on a sample level, MIDI files are typically much smaller than audio files. For instance, the .wav file of *I Saw Her Standing There* which we use in our data set has a size of 92.6 megabytes, while our MIDI files of the same song have file sizes ranging from

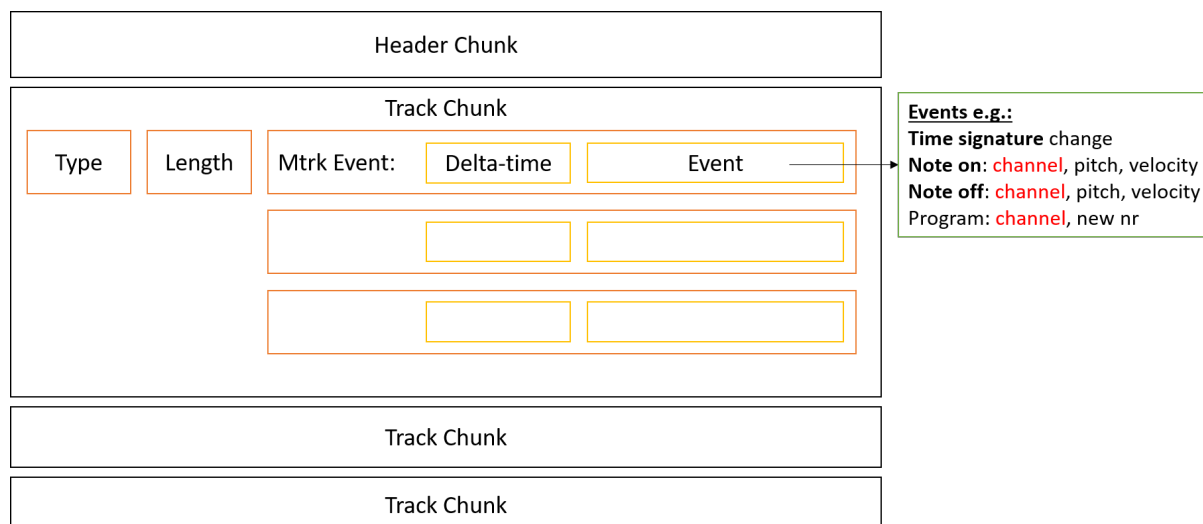


Figure 2.17: Structure of a MIDI file. A MIDI file has a header chunk and a number of track chunks. Each track chunk has a number of $\langle \text{delta-time}, \text{event} \rangle$ pairs. The most interesting events for our application are note on and note off events.

just 27.3 to 28.8 kilobytes. This storage efficiency hugely contributed to the popularity of MIDI files before the advent of the compressed file format MP3. Up until now, MIDI files are still used in resource-scarce settings such as karaoke machines. This has led to an abundance of MIDI files today. Recently, Raffel (2016) obtained as much as 178.561 MIDI files with unique MD5 checksums (i.e. a widely used 128-bit hash that is computed on a file) through a large-scale web scrape, resulting in the “Lakh MIDI Dataset”.

The structure of a MIDI file is illustrated in Figure 2.17. A MIDI file consists of a header chunk and a number of track chunks. The header chunk specifies the number of track chunks and the duration of a tick, which is the default time unit in a MIDI file. Each track chunk has a number of $\langle \text{delta-time}, \text{event} \rangle$ pairs. The delta-time is the time that has passed since the previous event, measured in ticks. An event is either a MIDI event, system exclusive (sysex) event or meta event. For a detailed description of these events and MIDI in general, we refer the reader to Guérin (2009). In this section, we will only consider a subset of events. The most common MIDI events in an average MIDI file are note on, note off and program change events. Note on and note off events respectively specify the start and end of a note and have a channel, pitch and velocity. Both pitch and velocity are integers between 0 and 127. A larger pitch value results in a higher sound: a pitch of zero corresponds to a C0, while the highest possible pitch (127) is the G10. If the velocity is high, a loud sound will be heard, while a velocity of zero only produces silence. A channel can be seen as a part in a full score: each channel is at any time mapped to a program number. These programs, or patches, determine the instrument sound. For example, program 1 is the acoustic grand piano, while a channel with program 59 will play the tuba for us. The program of a channel can be changed by program change events.

Based on note on and note off events, we can easily extract a piano-roll representation of a MIDI file. This is a time-frequency visualization that strongly reminds of piano rolls for pianola or reproducing piano. The piano-roll representation of a MIDI for the Beatles song *I Saw Her Standing There* is shown in Figure 2.18. When comparing this figure to

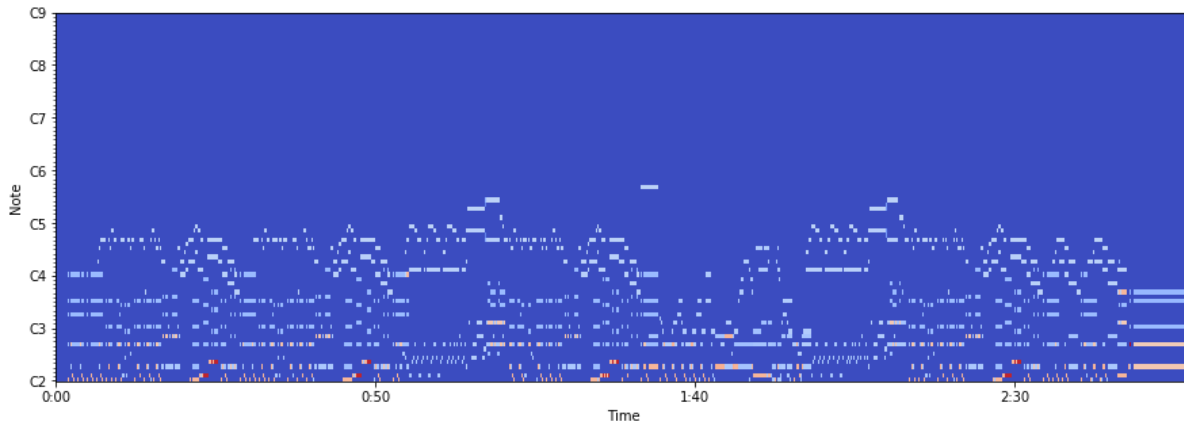


Figure 2.18: Piano-roll representation of *I Saw Her Standing There* by The Beatles. Note that this representation is comparable to the CQT-spectrogram, but way less noisy thanks to the absence of partials in MIDI.

the CQT-spectrogram in Figure 2.15, it becomes clear that the piano-roll representation is much more “clean”: thanks to the absence of partials, it is way easier to extract notes from MIDI than from audio.

Raffel and Ellis (2016a) describe the sources of information available in a MIDI file. First, MIDI files are naturally suited to be used as transcriptions of pieces of music, thanks to the way they are specified. At each position in the file, we can infer exactly which instrument plays which note. To extract this kind of information, various software libraries were developed, for example the MIDI Toolbox (Eerola and Toiviainen, 2004) and `pretty_midi` (Raffel and Ellis, 2014). Second, we can gather timing information from MIDI files, as there exist events for tempo changes and time signatures. Third, all 24 possible major and minor keys can be specified in a key change event. Fourth, lyrics can be added to MIDI transcriptions by the use of lyrics meta-events. Finally, software libraries like `jSymbolic` (McKay and Fujinaga, 2006) and `music21` (Cuthbert and Ariza, 2010) can be used to compute higher-level features.

To summarize, in this section we have seen that MIDI files are a symbolic music representation from which we can extract all kinds of musical information, including a note transcription, using one of the software libraries developed for this purpose. Thanks to their storage efficiency, MIDI files abound on the Internet. This makes MIDI files particularly interesting for ACE research.

2.5.3 Tab representation

Guitar tablatures and chord sheets are collectively known as “tabs”. In contrast to traditional musical scores, **(guitar) tablature** indicates the instrumental fingering rather than musical pitches. These tablatures are usually represented using an ASCII text notation, in which each line represents a string of the instrument. As reading tablature requires little musical training and tabs can be written and read without any specific software, they are very popular: millions of guitar tablature files can be found on web-

```

e|-----|-----|
B|-----|-----|
G|-----|-----|
D|-----|-----|
A|2--2--2--2--2--2--|2--2--2--2--2--2--|
E|0--0--0--0--0--0--|0--0--0--0--0--0--|
                                Well, she was just

          E7
e|-----0--0--0--0--0--0--|-----0--0--0--0--0--0--|
B|-----3--3--3--3--3--3--|-----3--3--3--3--3--3--|
G|-----1--1--1--1--1--1--|-----1--1--1--1--1--1--|
D|-----2--2--2--2--2--2--|-----2--2--2--2--2--2--|
A|-----2--2--2--2--2--2--|-----2--2--2--2--2--2--|
E|-----0--0--0--0--0--0--|-----0--0--0--0--0--0--|
                                seventeen,                                you know

          A7                                E7
e|-----3--3--3--3--3--3--|-----0--0--0--0--0--0--|
B|-----2--2--2--2--2--2--|-----3--3--3--3--3--3--|
G|-----2--2--2--2--2--2--|-----1--1--1--1--1--1--|
D|-----2--2--2--2--2--2--|-----2--2--2--2--2--2--|
A|-----0--0--0--0--0--0--|-----2--2--2--2--2--2--|
E|-----x--x--x--x--x--x--|-----0--0--0--0--0--0--|
                                what I mean                                and the

```

Figure 2.19: Guitar tablature excerpt

```

I SAW HER STANDING THERE
THE BEATLES

[Verse]
          E7                                A7                                E7
Well she was just seventeen and you know what I mean
                                B7
And the way she looked was way beyond compare
          E                                E7                                A7                                Am/C
So how could I dance with another oh,
          E7                                B7                                E7
when I saw her standing there

[Verse]
          E7                                A7                                E7
Well she looked at me and I, I could see
                                B7
That before too long I'd fall in love with her
          E                                E7                                A7
She wouldn't dance with another
          Am/C                                E7                                B7                                E7
Oh, when I saw her standing there

```

Figure 2.20: Chord sheet excerpt

sites like Ultimate Guitar¹ (Macrae and Dixon, 2011). Figure 2.19 shows an example of guitar tablature. Chords can be extracted from tabs very easily. For example, above the word “seventeen” we see the fingering combination $\langle 0, 3, 1, 2, 2, 0 \rangle$, which means that a chord is played using open E strings, the third fret on the B string, the first fret on the G string and the second fret on the D and A string. This way, the notes E4, D4, G#3, E3, B2 and E2 are played on the six strings. These notes form the seventh chord E7, as they consist of the notes E, G#, B and D. In many tablatures, including the example in Figure 2.19, chords are represented twice by also adding a chord. This makes it even more easy to extract chord information from guitar tablature. However, note that tablature does not contain any timing information, in contrast to audio and MIDI representations.

A **chord sheet** is a lyric sheet, in which chord symbols are placed above the lyric syllables with which they have to be timed. Chord sheets can also be very compactly represented in ASCII text notation. They can be found in abundance on the Internet. Figure 2.20 is a fragment from a chord sheet for *I Saw Her Standing There*. We see that chord information can be extracted directly from the chord sheet. However, similar to guitar tablature, there is no timing information available in chord sheets. In addition, it is quite common for chord sheets to represent only a single verse and chorus, as the chords of other verses and choruses in pop songs are often the same.

In this section, we learned that tabs consist of guitar tablatures and chord sheets. Both types of tabs are created by music enthusiasts and can be found in abundance on websites like Ultimate Guitar. However, as there are no restrictions on the authorships of tabs, many tabs are erroneous or incomplete. Therefore, it is not trivial to sift the wheat from the chaff and select only the high-quality tabs.

¹<https://www.ultimate-guitar.com/>

2.6 Summary of musical background

This chapter provided an introduction to the concepts of pitch, music notation, intervals, chords and music representations. We have seen that the **pitch** of musical tones refers to the degree of highness of sound and that the frequency is the number of vibrations in the sound wave. A tuning system discretizes the space of all different pitches. In this research project, we use twelve-tone equal-tempered tuning, in which the frequency ratio between of two subsequent scale steps equals $2^{\frac{1}{12}}$. A distance of twelve semitones is called an octave. We can group tones that are an octave apart in the same pitch class. In modern Western music notation, tones are given note names based on a pitch class name (i.e. one letter from {A, B, C, D, E, F, G}, possibly combined with an accidental) and an octave number. The reference note A4 has a frequency of 440 Hz. Music is **notated** on a staff. The pitch of each note is determined by its vertical position on the staff. **Intervals** are defined as the distance between two pitches and are based on the notion of semitone. Interval names consist of a number and a quality. **Chords** are groups of notes sounding at the same time, consisting of tones from at least different pitch classes. Most chords are built by stacking stacks of thirds on a root note. Triads are chords that consist of tones from three pitch classes and consist of the root, plus the third and fifth above it. Seventh chords, or tetrads, consist of the root, third, fifth and seventh. There exist multiple inversions for each chord, which are dependent on the relation of root note and bass note. We have seen that there exist various notations for chords. In this project, we use the chord grammar proposed by Harte et al. (2005) as this grammar is unambiguous, context-independent and easy to use.

In Section 2.5 we discussed the three different music representations which are used in this research project. The audio representation is an digitization of the waveform, obtained by sampling and quantization. Pitch information cannot be read from the audio directly, but can be estimated using Fourier analysis. MIDI files are another music representation which, thanks to their storage efficiency, abound on the Internet. Interestingly, MIDI files are defined in such a way that we can easily extract all kinds of musical information, including pitch information. The third and final music representation we study are tabs, which is an umbrella term for guitar tablature and chord sheets. Tabs can be found in abundance on the Internet. Chords can be extracted almost directly from tabs, although tabs do not contain timing information and the chord labels are not always reliable.

Chapter 3

Framework of proposed system “DECIBEL”

In this chapter, I present the framework of DECIBEL: the proposed system for the Detection of Chords Improved By Exploiting Linked symbolic formats. In Chapter 1 we have seen that existing chord estimation techniques that are only based on audio have some limitations: the performance of the MIREX ACE submissions seems to have reached a glass ceiling and some modern systems suffer from overfitting to subjective reference annotations. In Section 2.5 we saw that the symbolic representations MIDI and tab have the convenient property that it is very easy to extract notes and chords from them. DECIBEL exploits this property by aligning MIDIs and tabs to the corresponding audio file, extracting chord sequences from each of the representations and using data fusion to combine the resulting chord sequences. DECIBEL’s framework is summarized in Figure 3.1 and will be explained concisely in this section.

3.1 DECIBEL’s framework in a nutshell

The DECIBEL system has a data set of audio, MIDI files and tabs at its disposal. MIDI files and tabs are obtained by a **web scrape**. These files are manually matched, based on meta-data. The dataset and matching process are described in Section 3.2.

For each song, each of the three representations (audio, MIDI and tabs) is mapped to an **audio-timed chord sequence**, which is a sequence of chord events. Chord events are 3-tuples, consisting of a start time, end time and chord label. The possible chord label values are specified by the chosen chord vocabulary. The method for this chord estimation step depends on the representation: we used three different methods for audio, MIDI and tab representations, as specified below:

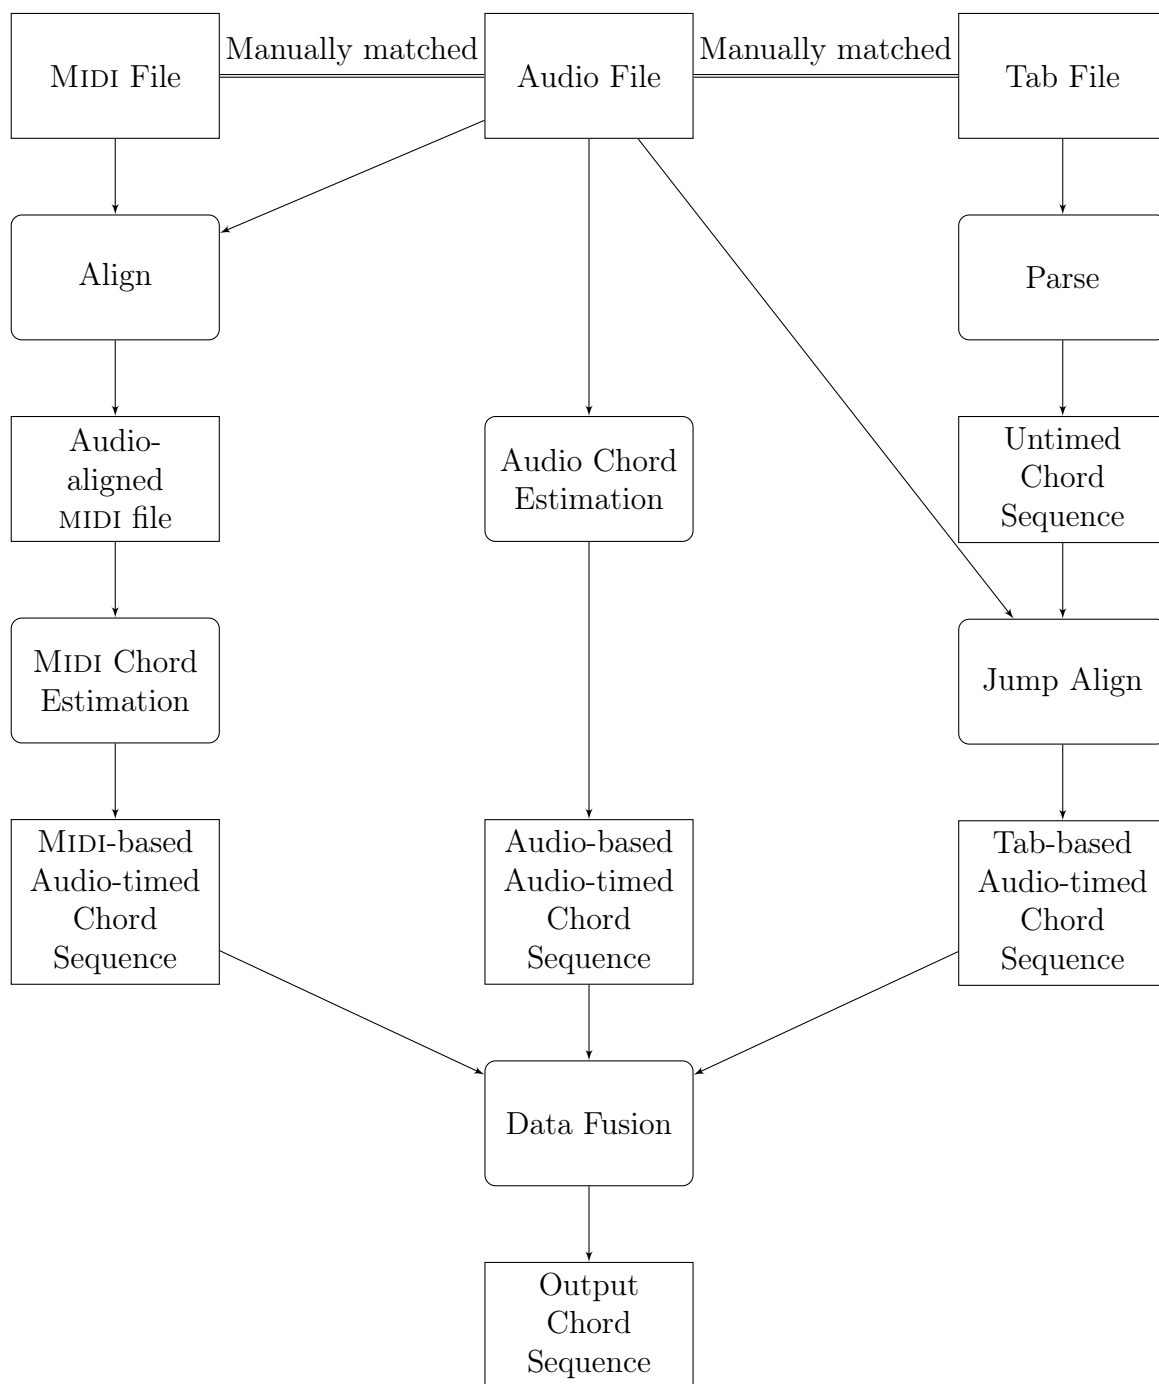


Figure 3.1: Diagram of DECIBEL

- Audio:** DECIBEL estimates chords from audio data using existing **audio ACE** techniques.
- MIDI:** In order to estimate chords from MIDI, DECIBEL first aligns each MIDI file to the audio recording using **audio-symbolic alignment** techniques. Then, chords are extracted from the audio-aligned MIDI file using a pattern-matching technique for **chord estimation in symbolic music**. This way, DECIBEL obtains the chord sequences with the correct start and end times for the original audio file.
- Tabs:** For the tabs, DECIBEL can easily find the chord labels by **parsing** the ASCII text. Consequently, the system aligns them to the audio using **Jump Alignment**.

Each of the three procedures for representation-dependent extraction of audio-timed chord sequences can be considered as a subsystem of DECIBEL. In Chapters 4, 5 and 6, I will describe the audio, MIDI and tabs subsystems in detail.

At this point, we have a rich harmonic representation, consisting of possible chord sequences for the song, obtained from symbolic and audio representations. As a final step, I use **data fusion** to combine these chord sequences into one final chord sequence. The data fusion method is treated in Chapter 7.

In the remainder of this section, I describe the collection of the data set in detail in Section 3.2. Also, I describe the performance evaluation measures, which are used to test both the representation-specific subsystems as well as the data fused result, in Section 3.3.

3.2 Collection of data set

DECIBEL uses a data set of audio, MIDI files and tabs. This data set is based on a subset of the Isophonics Reference Annotations (Mauch et al., 2009). The Isophonics data set contains chord annotations for 180 Beatles songs, 20 songs by Queen, 7 songs by Carole King and 18 songs by Zweieck. In this project, I only use the songs by the Beatles and Queen, as there were no MIDI or tabs for Zweieck available and there were some inconsistencies in the Carole King annotations. Using this 200 song data set has three advantages:

- The music by The Beatles and Queen is popular music and therefore the correct genre for our dataset;
- Thanks to the popularity of the two bands, it is easy to find MIDI and tab files for the songs in the data set;
- The chord labels have been carefully checked and have been used for many years by the MIR community.

I used the audio as provided on the CDs in Table 3.1. A complete list of all song names and the index I assigned to them, is given in Appendix A. After collecting the audio

Artist	Album	ID	# songs
The Beatles	Please Please Me	CDP 7 46435 2	14
The Beatles	With the Beatles	CDP 7 46436 2	14
The Beatles	A Hard Day’s Night	CDP 7 46437 2	13
The Beatles	Beatles For Sale	CDP 7 46438 2	14
The Beatles	Help!	CDP 7 46439 2	14
The Beatles	Rubber Soul	CDP 7 46440 2	14
The Beatles	Revolver	CDP 7 46441 2	14
The Beatles	Sgt. Pepper’s Lonely Hearts Club Band	CDP 7 46442 2	13
The Beatles	Magical Mystery Tour	CDP 7 48062 2	11
The Beatles	The Beatles (the white album)	CDS 7 46443 8	30
The Beatles	Abbey Road	CDP 7 46446 2	17
The Beatles	Let It Be	CDP 7 46447 2	12
Queen	Greatest Hits I	Parlophone, 0777 7 8950424	14
Queen	Greatest Hits II	Parlophone, CDP 7979712	6

Table 3.1: Isophonics Reference Annotations

files and annotations, we need to find MIDI and tab files and match them to the songs in our data set. First, I searched on the Internet for MIDI files of the aforementioned 200 songs. I downloaded MIDI files from 9 websites¹²³⁴⁵⁶⁷⁸⁹. This way, I found 770 MIDI files with unique MD5-checksums, so multiple MIDI files (3.85 on average) map to a single audio file. I matched the MIDI and audio files by hand, based on the MIDI file name. Furthermore, I obtained tabs from Ultimate Guitar¹⁰. I first automatically scraped all tabs from The Beatles and Queen from Ultimate Guitar’s website. Then, I matched the tabs to the audio files by hand, based on song title. Tabs from songs that were not in the dataset, were discarded. This resulted in 1668 matched tabs (974 chords and 694 guitar tablature files).

Some statistics on the data set are shown in Figure 3.2. A typical song in our data set has a duration of 2 to 3 minutes, consists of 50-100 chord segments and is matched to 3 to 5 MIDI files and 5 to 10 tabs.

¹<http://beatlesnumber9.com>

²<http://bmh.webzdarma.cz>

³<http://davidbmidi.com>

⁴<http://earlybeatles.com/midi>

⁵<http://en.midimelody.ru>

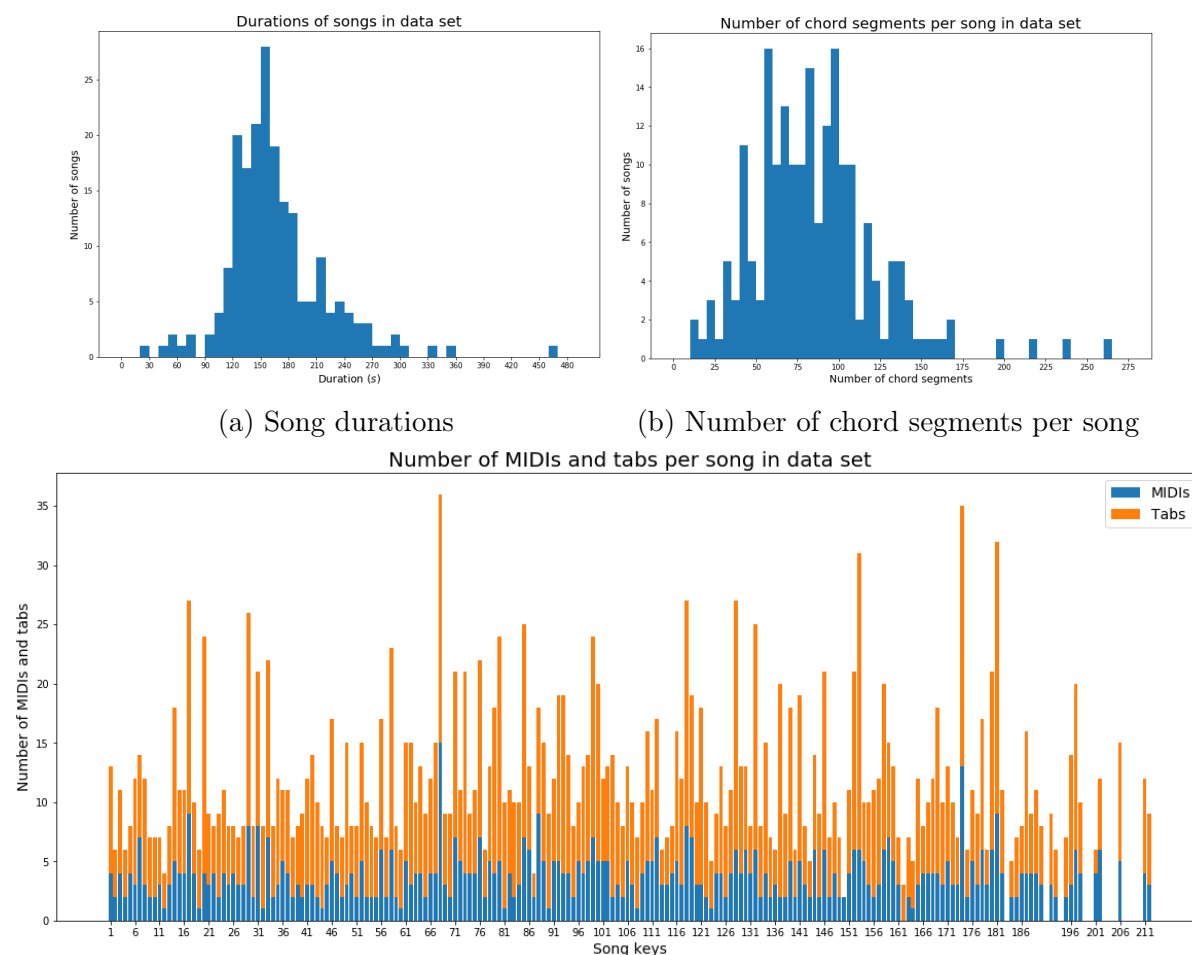
⁶<http://queen.wz.cz/midi>

⁷<http://www.angelfire.com>

⁸<http://www.dongrays.com>

⁹<http://www.rppmf.com>

¹⁰<https://www.ultimate-guitar.com>



(c) Number of MIDIs and tabs per song. Note that some songs (e.g. song number 208) are not matched to any MIDI or tab file. These are songs on the Queen CD's for which there were no reference annotations in the Isophonics data set.

Figure 3.2: Data set statistics. A typical song in our data set has a duration of 2 to 3 minutes, consists of 50-100 chord segments and is matched to 3 to 5 MIDI files and 5 to 10 tabs.

3.3 Performance evaluation

In order to evaluate both the performance of DECIBEL’s representation-specific subsystems and its final output chord sequence, we need evaluation measures. The quality of a chord sequence is usually determined by comparing it to a ground truth created by one or more human annotators. Commonly used **data sets** with chord annotations, which are also used in the MIREX ACE contest, are Isophonics¹¹, Billboard¹², RobbieWilliams¹³, RWC-Popular¹⁴ and USPOP2002Chords¹⁵. As stated before, DECIBEL uses the Isophonics data set, augmented with matched MIDI and tab files.

The standard quality measure to evaluate the quality of an automatic transcription is **chord symbol recall (CSR)** (Harte, 2010). This measure is also used in the MIREX ACE contest¹⁶. CSR is the summed duration of time periods where the correct chord has been identified, normalized by the total duration of the song. Until 2013, MIREX used an approximate, frame-based CSR calculated by sampling both the ground-truth and the automatic annotations every 10 ms and dividing the number of correctly annotated samples by the total number of samples. Since 2013, MIREX has used segment-based CSR, which is more precise and computationally more efficient. The formula for segment-based CSR is given in Equation 3.1. We consider the ground-truth annotation A as a sequence of segments S_A and the estimated annotation E as a sequence of segments S_E . The duration of a segment is notated as $|\cdot|$.

$$CSR_T(S_E, S_A) = \frac{\sum_{S_A^j} \sum_{S_E^i} |S_E^i \cap S_A^j| \cdot \mathbb{M}_T(S_A^j, S_E^i)}{\sum_{S_A^j} |S_A^j|} \quad (3.1)$$

\mathbb{M}_T is a matching function as defined by Equation 3.2, in which T denotes the comparison method used to evaluate the result of the matching function.

$$\mathbb{M}_T = \begin{cases} 1 & \text{if } X \text{ matches } Y \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

The comparison method is dependent of the **chord vocabulary**: both the chord labels ground truth annotations and the estimated annotations are mapped to a limited set of chord labels, as specified by the chord vocabulary. MIREX uses the following five chord vocabularies:

¹¹<http://isophonics.net/datasets>

¹²<http://ddmal.music.mcgill.ca/research/billboard>

¹³https://www.researchgate.net/publication/260399240_Chord_and_Harmony_annotations_of_the_first_five_albums_by_Robbie_Williams

¹⁴<https://github.com/tmc323/Chord-Annotations>

¹⁵<https://labrosa.ee.columbia.edu/projects/musicsim/usp2002.html>

¹⁶http://www.music-ir.org/mirex/wiki/2017:Audio_Chord_Estimation

Ground Truth	C					F			G		C			
Annotation A	C	B	C	F#	C	F	B	F	B	G	C			
Annotation B			Am			F			G		C			
$t =$	0	1	2	3	4	5	6	7	8	9	10	11	12	13

Figure 3.3: Annotation A has a CSR of $9/13 = 69.2\%$ and Annotation B has a CSR of $8/13 = 61.5\%$. However, Annotation B is the preferred chord sequence.

1. Only the chord root note (C, D, ..., B), or no-chord (N);
2. Major and minor: {N, maj, min};
3. Seventh chords: {N, maj, min, maj7, min7, 7};
4. Major and minor with inversions: {N, maj, min, maj/3, min/b3, maj/5, min/5};
5. Seventh chords with inversions: {N, maj, min, maj7, min7, 7, maj/3, min/b3, maj7/3, min7/b3, 7/3, maj/5, min/5, maj7/5, min7/5, 7/5, maj7/7, min7/b7, 7/b7}.

Two chords match if and only if they are mapped to exactly the same chord label in the vocabulary. For example, the mapping to chord vocabulary 1 only preserves the root note. The chords C:maj and C:min will both be mapped to a C chord, so these chords would match in the first chord vocabulary. However, in all other vocabularies they would be mapped to C:maj and C:min, which do not match.

For results that are calculated for the whole data set, we weigh the CSR by the length of the song when computing an average for a given corpus. This final number is referred to as the **weighted chord symbol recall (WCSR)**. Calculating the WCSR is basically the same as treating the data set as one big audio file, and calculating the CSR between the concatenation of all ground-truth annotations and the concatenation of all estimated annotations.

The CSR correctly indicates the accuracy of an ACE algorithm in terms of whether the estimated chord for a given instant in the audio is correct. It is therefore widely used in the evaluation of ACE systems. However, the annotation with the highest CSR is not always the annotation that would be considered the best by human listeners. As an example, examine Figure 3.3. Here we see two estimated annotations. Although Annotation A has the higher CSR, most musicians would prefer Annotation B: Annotation A is clearly over-segmented, which makes it difficult to play and unnatural to listen to. On the other hand, human listeners would consider the 5 seconds with the wrong chord in Annotation B as just one major mistake.

Just measuring the (weighted) chord symbol recall is therefore not enough: we also need a metric for chord segmentation quality. For this purpose, Mauch (2010) proposed the use

of the directional hamming distance $h(S, S^0)$. It describes how fragmented segmentation S is with respect to segmentation S^0 , according to Equation 3.3:

$$h(S, S^0) = \sum_{i=1}^{N_{S^0}} (|S_i^0| - \max_j |S_i^0 \cap S_j|) \quad (3.3)$$

In our example, let Annotation A be S^A , let Annotation B be S^B and let the ground truth be S^G . Then, $h(S^A, S^G)$ is a measure of over-segmentation of S^A with respect to S^G . Indeed,

$$h(S^A, S^G) = \sum_{i=1}^{N_{S^G}} (|S_i^G| - \max_j |S_i^G \cap S_j^A|) = (5 - 1) + (3 - 1) + (2 - 1) + (3 - 3) = 7$$

This is a high value, so the directional Hamming distance shows that S^A is over-segmented. A high Hamming distance in the opposite direction indicates under-segmentation.

$$h(S^G, S^A) = \sum_{i=1}^{N_{S^A}} (|S_i^A| - \max_j |S_i^A \cap S_j^G|) = 13 * (1 - 1) = 0$$

so S^A is not under-segmented. The segmentation of Annotation B is identical to the ground truth:

$$h(S^B, S^G) = h(S^G, S^B) = 0$$

We can easily transform the directional Hamming distance into a quality measures for over-segmentation and under-segmentation using the following equations:

$$\text{OverSegmentation}(S_E, S_A) = 1 - \frac{h(S_E, S_A)}{\sum_{i=1}^{N_{S_A}} |S_{A_i}|} \in [0, 1] \quad (3.4)$$

$$\text{UnderSegmentation}(S_E, S_A) = 1 - \frac{h(S_A, S_E)}{\sum_{i=1}^{N_{S_A}} |S_{A_i}|} \in [0, 1] \quad (3.5)$$

OverSegmentation or UnderSegmentation values near 0 correspond to highly under- and over-segmented annotations respectively. High values indicate good segmentation quality. Equations 3.4 and 3.5 can be combined into an all-in-one segmentation measure using Equation 3.6. We use the minimum, so an annotation only gets a high segmentation quality if neither under-segmentation nor over-segmentation is dominant.

$$\text{Segmentation}(S_E, S_A) = \min \left\{ \begin{array}{l} \text{OverSegmentation}(S_E, S_A) \\ \text{UnderSegmentation}(S_E, S_A) \end{array} \right. \quad (3.6)$$

Chapter 4

Automatic Chord Estimation on audio

In this chapter, we focus on the subsystem of DECIBEL that extracts chord sequences from the audio representation, which is schematically summarized in Figure 4.1. There exist multiple methods to extract chord labels from audio data. In Section 4.1, I give an overview of existing methods for ACE in audio, showing a variety in implementations. Seven state-of-the-art systems are used in the audio ACE subsystem of DECIBEL, as described in Section 4.2. Finally, I evaluate these seven systems on our 200 song data set in Section 4.3 and conclude in Section 4.4.




Figure 4.1: Diagram of DECIBEL’s audio subsystem

4.1 Related work on audio ACE

This section gives an overview of existing methods for ACE on audio. Fujishima (1999) was the first who considered ACE as a task on its own. Since then, a lot of researchers have buckled down to this subject.

Most methods use the following pipeline (McVicar et al., 2014): first, audio data is partitioned into a training set and a test set and features are calculated on both partitions. Subsequently, the features from the training set are used to train the parameters of a model. The chord labels for the test set are then estimated using this trained model. Finally, the performance of the system is evaluated by comparing the labels calculated by the model to the reference (ground truth) chord labels. Feature extraction and models are described in Sections 4.1.1 and 4.1.2 respectively. The standard pipeline is illustrated in Figure 4.2.

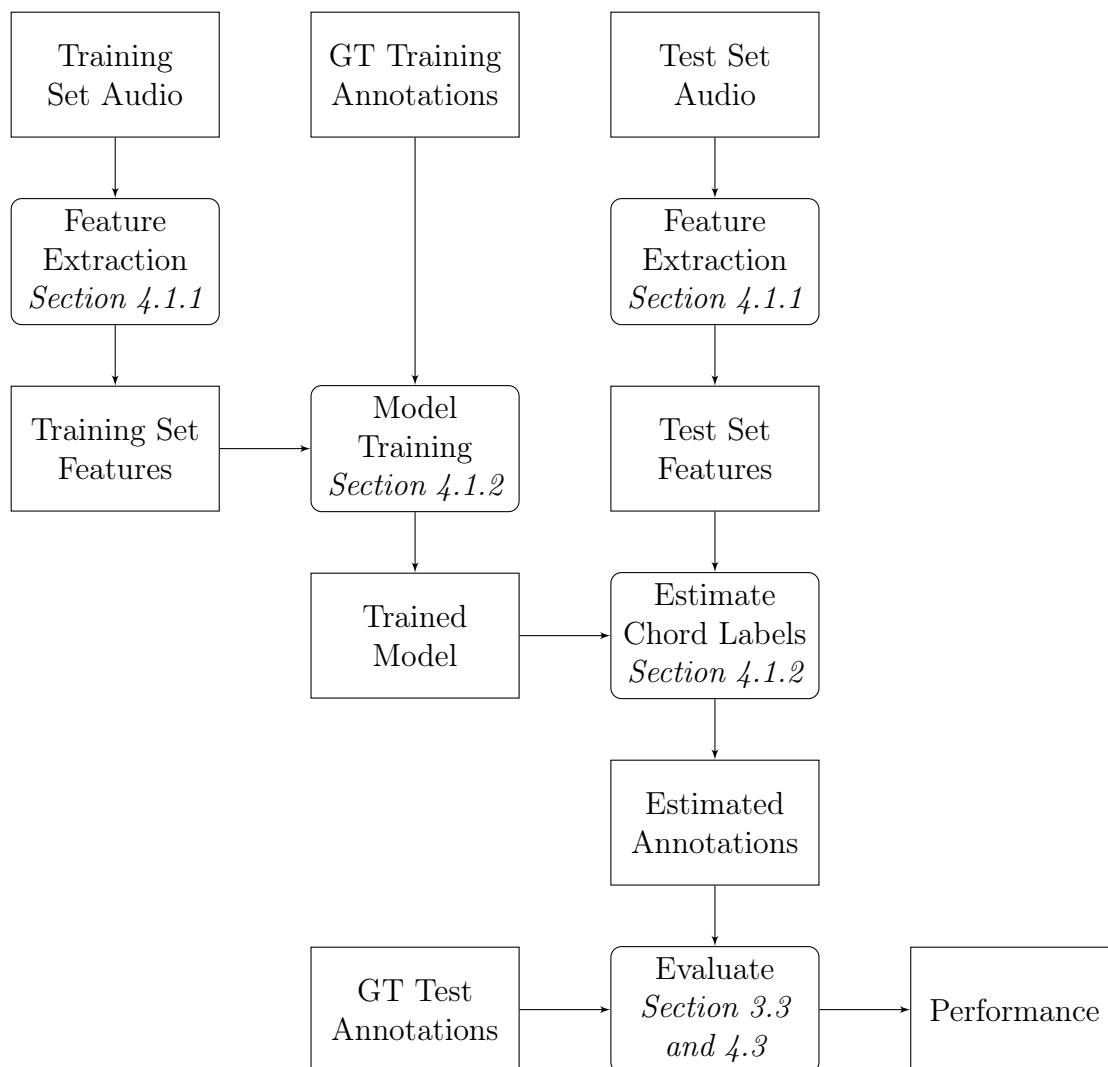


Figure 4.2: A prototypical audio ACE pipeline, which estimates and evaluates chord label sequences based on a data set of audio files. There exist multiple methods for feature extraction and modeling. These are described in Sections 4.1.1 and 4.1.2.

4.1.1 Feature extraction for audio ACE

The first step in audio chord estimation is feature extraction. Features are extracted for both the audio in the training set and the audio in the test set, as shown in Figure 4.2.

The first operation most ACE methods perform in feature extraction, is to calculate (a variation of) the chromagram (Wakefield, 1999). As we have seen in Section 2.5.1, the chromagram describes the pitch class salience over the duration of the audio, i.e. it represents which notes are present in each audio frame. For calculating the chromagram, the audio is first transformed to the **frequency domain** using either the Short Time Fourier Transform (STFT) or Constant-Q transform.

In many ACE systems, the spectrogram is **preprocessed** before calculating the chromagram. For example, Mauch (2010) removes background noise by median filtering of the spectrogram; Reed et al. (2009) remove the percussive elements of the spectrum and Pauws (2004) removes harmonics. As some popular music songs are not tuned at the standard pitch of $A4 = 440$ Hz, many methods compensate for tuning issues by computing the spectrogram at a multiple of the required frequency resolution: Sheh and Ellis (2003) doubled the frequency resolution of the spectrogram and Harte and Sandler (2005) even used the triple frequency resolution. As a next step, some authors remap the spectrogram, so it represents the human perception of pitch saliences more closely. Ni et al. (2012) for example use A-weighting, i.e. a technique to express the loudness as perceived by the human ear, to calculate loudness-based chromagrams. Finally, almost all methods discard octave information, as this is not relevant in chord estimation: the pitch A4 has the same function in a chord as the pitch A3 or A5.

Given the chromagram, some **post-processing** steps may help to prevent frequent chord changes in the predicted chord sequence. Fujishima (1999) used smoothing techniques as post-processing, while Bello and Pickens (2005) introduce beat-synchronous chromagrams, in which the chromagram is averaged between beat segments.

Some methods however use **other representations** than traditional chromagrams: Harte et al. (2006) introduce the Tonal Centroid feature; Mauch and Dixon (2008) calculate a distinct bass chromagram and in more recent work, Wu et al. (2017) calculate a 36-dimensional binary acoustic feature using a Deep Residual Network, trained on MIDI data. Sigtia et al. (2015) use Deep Neural Networks and Korzeniowski and Widmer (2016a,b) train a Convolutional Neural Network to automatically learn musically interpretable features from the spectrogram. Müller and Ewert (2010) introduce chroma DCT-reduced log pitch (CRP) features, which improve timbre invariance. These CRP features are implemented in the ACE system by O’Hanlon et al. (2017).

In this subsection, we saw various ways of extracting features from audio. These features are extracted from the audio of both the training set and the test set, and serve as an input for a model, as we will see in the next session.

4.1.2 Audio ACE models

As a next step, we need to estimate chord labels, given the feature vectors of the test set. Although early methods for calculating the chord labels from feature vectors rely on template matching, recent works typically use statistical models or neural networks. Usually, these models are first trained on the training set, which is a subset of the data set for which ground truth chord labels are known. The model is then evaluated on the other partition of the data set: the test set. This section briefly indicates the variety of models used in audio ACE.

Sheh and Ellis (2003) were the first to use the **Hidden Markov Model (HMM)** and Expectation-Maximization algorithm to train a model for ACE, in which the chord labels are represented by hidden states and the features correspond to the observed states. HMMs model the joint probability distribution $P(\mathbf{X}, \mathbf{y})$ over the feature vectors X and the chord labels y . We will have a more detailed look at HMMs in Section 6.2.2, as this type of model is also used in the Jump Alignment algorithm (with which DECIBEL aligns tabs to audio). Some methods, for example the system proposed by Shenoy and Wang (2005), incorporate key information in a two-chain HMM. Scholz et al. (2009) introduce higher-order HMMs, in order to better model the complexity of music.

An alternative model is the **Dynamic Bayesian Model (DBM)**, which is a Bayesian network that relates variables to each other over adjacent time steps. It was introduced for ACE by Mauch and Dixon (2010). This model has hidden nodes for metrical position, chord, bass note and musical key and observed nodes for bass and treble chromagrams. Thanks to its design, the DBM is able to integrate multiple pieces of musical context in a single model.

Burgoyne et al. (2007) used different variations of **Conditional Random Fields (CRF)**. CRF is a discriminative model: it models $P(\mathbf{y}|\mathbf{X})$, which is the conditional distribution of chord labels given a sequence of feature vectors. Linear-chain CRFs were also used in the system of Korzeniowski and Widmer (2016b).

In recent work, we see a trend towards using **Artificial Neural Network** models. This class of models is inspired by biological neural networks in our brains: they consist of a collection of connected nodes or neurons, often arranged in layers. Each of the neurons processes a signal that it receives from neurons of the previous layer, calculates some non-linear function on the sum of these inputs and passes the result on to the neurons in the next layer. The neurons and edges between them typically have weights that are adjusted in the training phase. This way, neural networks can learn a function that maps feature vectors to chord labels without modeling a probability distribution - provided that there is sufficient training data. Neural networks come in many flavors. For example, Sigтия et al. (2015) use a hybrid Recurrent Neural Network (RNN) model and Wu et al. (2017) combine a Bi-direction LSTM Network and CRFs.

4.2 Selected systems

In DECIBEL, we experiment with seven different audio ACE systems: the six submissions for the MIREX 2017 ACE competition, together with the Chordify algorithm. I selected these seven systems as they are state of the art. It is therefore, in contrast to inferior audio ACE systems, not trivial that we can improve these systems by incorporating information from MIDI files and tabs. The implementation of seven selected systems is summarized below.

- CHF:** The exact details of the Chordify algorithm are not public, but the current implementation is based on recent papers by McFee and Bello (2017) and Koops et al. (2017). CHF uses a deep convolutional-recurrent model for automatic chord recognition from the CQT spectrogram.
- CM2:** CM2 is the algorithm implemented in the Chordino plugin in Sonic Visualiser (Cannam et al., 2017). First, NNLS chroma features are extracted from the CQT spectrogram. Then, a fixed dictionary of chord profiles is used to calculate frame-wise chord similarities. Finally, the chord labels are computed using a HMM.
- JLW1, JLW2:** Both JLW algorithms are based on the random forest model, which is a collection of decision trees (Jiang et al., 2017). First, the signal is separated into a harmonic and percussive signal, using HPSS. Then the harmonic part is transformed to the frequency domain and NNLS chroma features are calculated on the result. Consequently, a random forest model is trained on these features. This results in a chordogram, which is a matrix that represents the emission probability for chord c at frame n . As a final step, the result is smoothed, and this is where versions JLW1 and JLW2 differ. JLW1 uses a HMM, similar to the one used by Chordino. JLW2 uses a beat tracking CRF model that only allows chord changes on beats or at half-beat positions.
- KBK1, KBK2:** Each of the two KBK algorithms is based on neural networks and only recognizes chords in the Major/Minor alphabet. KBK1 uses a deep neural network to extract chroma vectors, as shown in (Korzeniowski and Widmer, 2016a), while KBK2 learns features automatically by a fully convolutional neural network (Korzeniowski and Widmer, 2016b). Then, in both versions, the chord sequence is decoded using a linear-chain Conditional Random Field (CRF) as described in Korzeniowski and Widmer (2016b) and chords segments are aligned to beats using a beat tracker.
- WL1:** In WL1, the acoustic features are first calculated from the spectrogram of each music signal with a deep residual network. Then, the feature vectors are fed as a sequence into a Bidirectional Long Short Term Memory network, and a class likelihood vector is calculated for each frame. Finally, the class likelihood sequence is passed on to the trained CRF layer to decode the optimal label sequence.

4.3 Evaluation of audio ACE

In this section, I evaluate the seven audio ACE systems, of which I described the implementation in the previous section, on my data set¹. Figure 4.3 shows the performance of all seven systems in terms of CSR and segmentation measures. From the figure we see that the CSR of all seven systems is quite high, with a median value of at least 80%. KBK2 is the best performing system, but given these high CSR scores we can already conclude that all seven systems are state of the art. We also observe some outliers for all systems, which are songs for which the CSR is very low. Some outliers are song-specific, that is: (almost) all audio ACE systems perform badly on this song. Other outliers are specific for an ACE system.

First, we will look at song-specific outliers, which have low CSRs due to problems in the audio recording or in the ground truth labels. Some songs have **tuning issues**. This is for example the case in the song *Lovely Rita* from the Beatles. The song was originally performed in E major, but during mixdown the tape machine ran at a lower frequency, resulting in a pitch drop of a quarter tone (Lewisohn, 1989). Without this background information, it would therefore be very hard to decide if the song is in E major or in Eb major. From Figure 4.4 it becomes clear that all seven audio ACE systems choose the Eb major key: the chord labels of the ACE systems are quite consistently one semitone lower than the ground truth chord labels. This results in very low CSRs for all seven audio ACE systems, ranging from 0.2% (JLW2) to 20% (KBK2). The Beatles song *Wild Honey Pie* has tuning issues as well, resulting in CSRs from only 0.5% (JLW1) to 41% (KBK2). *Wild Honey Pie* is a short, experimental song, characterized by Indian influences, played with a lot of vibrato. The Queen song *Another One Bites The Dust* has some tuning issues too, albeit to a lower extent than *Lovely Rita* and *Wild Honey Pie*.

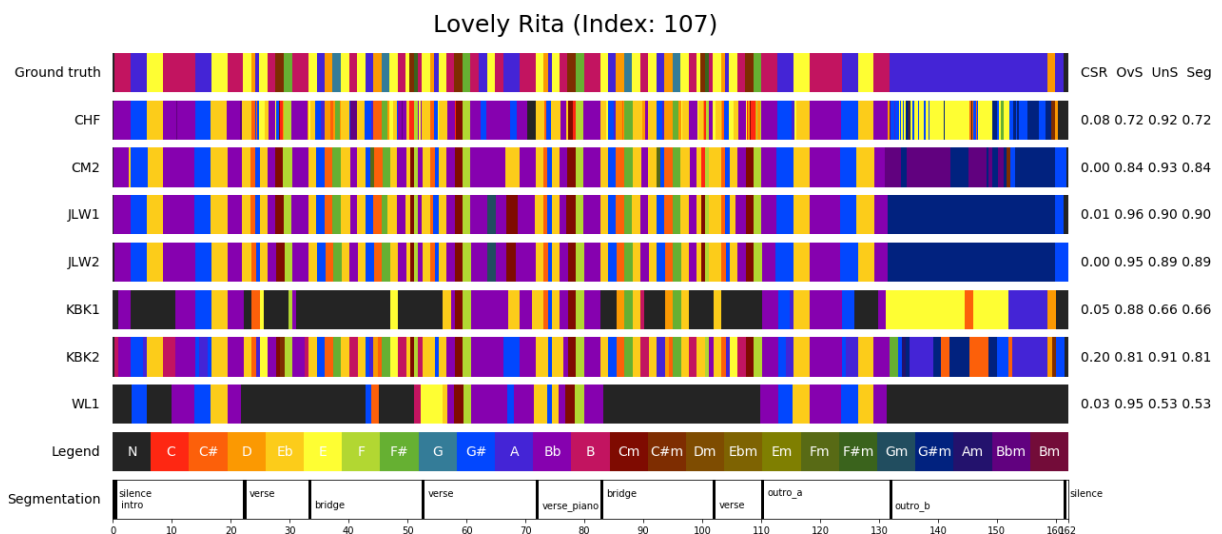




Figure 4.4: Chord sequences from seven different audio ACE systems on the Beatles song *Lovely Rita*, which has tuning issues. CSR = Chord Symbol Recall; OvS = oversegmentation; UnS = undersegmentation; and Seg = segmentation.


¹Note that the results are not exactly the same as the MIREX evaluation on the Isophonics data set, as we use a subset of this data set, consisting of the Beatles and Queen songs.



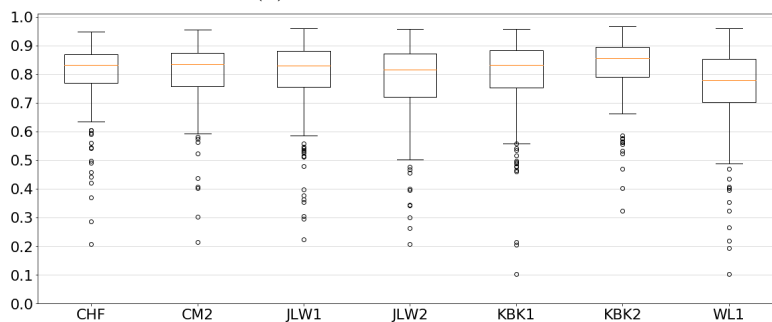
(a) Chord Symbol Recall



(b) Oversegmentation



(c) Undersegmentation



(d) Segmentation

Figure 4.3: Comparison of performance measures for Chordify (CHF) and MIREX 2017 ACE submissions. For a description of these performance measures, see Section 3.3. KBK2 is the best performing system in terms of CSR. CHF and CM2 tend to oversegment the song. On the other hand, WL1 inclines to undersegmentation.

In two other songs (*The Continuing Story of Buffalo Bill* and *Don't Pass Me By*), the chord labels of the six MIREX chord ACE systems are **shifted** in time compared to the ground truth data, as shown in Figure 4.5. This is probably due to the use of different audio versions in chord annotation and evaluation.




Figure 4.5: Chord sequences from seven different audio ACE systems on the Beatles song *Don't Pass Me By*, in which the chord labels are shifted in time. CSR = Chord Symbol Recall; OvS = oversegmentation; UnS = undersegmentation; and Seg = segmentation.

The Beatles song *Revolution* has another song-specific issue. This song is a sound collage: an experimental recording that is glued together from sound fragments, many of which are **non-harmonic**. Accordingly, we see in Figure 4.6 that a large part of the ground truth is represented with black, corresponding to the no-chord symbol. However, most audio ACE systems do not correctly recognize non-harmonic content as they are trained on tonal music.

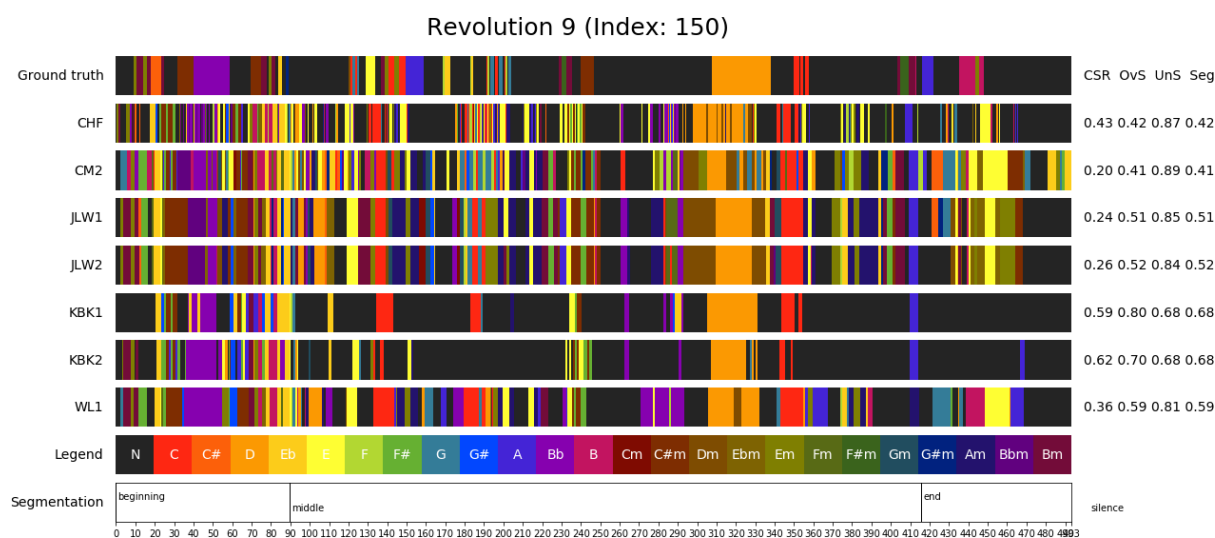


Figure 4.6: *Revolution 9* is a sound collage, consisting of many non-harmonic sounds. CSR = Chord Symbol Recall; OvS = oversegmentation; UnS = undersegmentation; and Seg = segmentation.

Apart from song-specific problems (like tuning issues, shifted recordings and non-harmonic notes), there are also audio ACE-specific issues. From Figure 4.3a, we can see that KBK2 performs the best, while CM2 performs the worst in terms of CSR. CHF and CM2 tend to oversegment the song, considering the relatively low values for oversegmentation in Figure 4.3b. On the other hand, WL1 inclines to undersegmentation. We see this for example in Figure 4.7: in this example we see that CHF and CM2 create too many segments and correspondingly get relatively low oversegmentation scores. On the other hand, WL1 creates too few segments, resulting in undersegmentation: WL1 here has a undersegmentation score of only 0.53.




Figure 4.7: Chord sequences from seven different audio ACE systems on the Beatles song *Another Girl*. CSR = Chord Symbol Recall; OvS = oversegmentation; UnS = undersegmentation; and Seg = segmentation.

4.4 Conclusion

In this chapter, we have described the implementation and performance characteristics of seven systems for audio ACE. By evaluating the audio ACE systems on our 200 song data set, we learned that all seven systems are state of the art, as they have high chord symbol recalls. However, different ACE systems behave differently: some systems tend to undersegment the chord sequence, while others tend to oversegmentation. In my research, I will use of each of these systems: I will compare the performance of the original audio ACE system to the data fused result, which incorporates not only the audio ACE system but also information from MIDI and tab files.

Chapter 5

Automatic Chord Estimation on MIDI

In this chapter, I discuss the subsystem of DECIBEL that detects an audio-timed chord sequence based on a MIDI file. This subsystem is illustrated below, in Figure 5.1. In order to receive audio-timed chord labels from a MIDI file, DECIBEL first finds an optimal alignment from the MIDI file to the audio file, realigns the MIDI file using this alignment and then uses a MIDI chord recognizer to estimate the chord labels on the realigned MIDI file. In Section 5.1 I will explain the alignment step and in Section 5.2 the MIDI chord recognition step. We then have audio-timed chord labels for each MIDI file. However, there may be MIDI files that are unsuitable to use, because they are bad transcriptions (i.e. the notes in the MIDI file are not consistent with the notes in the audio file), because they are in a wrong transposition or because they are a transcription of just one part of the song. It may help to exclude these MIDI files from the data fusion step. We experiment with MIDI selection methods in Section 5.3 and conclude in Section 5.4.

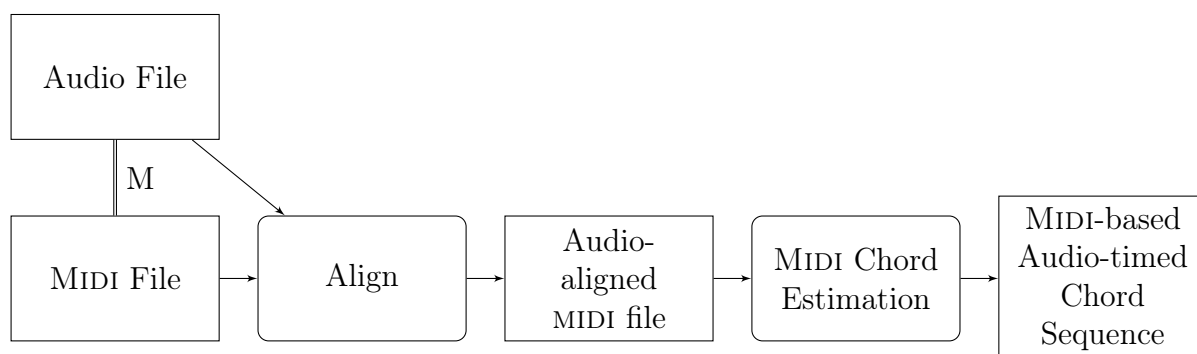


Figure 5.1: Diagram of DECIBEL’s MIDI subsystem; the M between Audio File and MIDI file indicates that they are manually matched.

5.1 MIDI-to-audio alignment

Music alignment is the procedure which, given any position in one representation of a piece of music, determines the corresponding position within another representation. It is also called *music synchronization*. In the MIDI subsystem, DECIBEL aligns each of the MIDI representations to the corresponding audio representation. This section explains the MIDI-to-audio alignment procedure. MIDI-to-audio alignment is an important task in MIR: a lot of research has already been conducted on this subject. Therefore, we will first study related work on alignment in Section 5.1.1. We can then make a well-grounded choice for the alignment system implemented in DECIBEL. This alignment system is described in Section 5.1.2 and evaluated in Section 5.1.3.

5.1.1 Related work on MIDI-to-audio alignment

There are two variants of MIDI-to-audio alignment: on-line and off-line alignment. In **off-line alignment**, the full performance is accessible for the alignment process. By contrast, in **on-line alignment**, the aligner processes the data in real-time as the signal is acquired. On-line alignment is also known as score following and has applications such as automatic accompaniment, audio editing and automatic turning of score pages. In general, off-line techniques work better than on-line systems, because they do not need to calculate the alignment in real-time and they can access the full performance at all times. For our project, the full performance is accessible from the beginning of the alignment process and the alignment does not have to be calculated in real-time. Therefore, we use off-line audio-symbolic alignment techniques.

The alignment method is typically based on either **statistical approaches** (for example HMMs, as used by e.g. Cuvillier and Cont (2014)) or **Dynamic Time Warping** (Carabias-Orti et al., 2015; Raffel and Ellis, 2016b; Arzt, 2016; Lajugie et al., 2016). In this project, we focus on Dynamic Time Warping (DTW) approaches, as the DTW algorithm is efficient, is conceptually simple and calculates an easily interpretable alignment confidence score.

In the remainder of this section, I first give a general introduction to DTW. Subsequently, we see how DTW can be applied to alignment of music.

An introduction to Dynamic Time Warping

In DTW (Sakoe and Chiba, 1978) we try to find the optimal alignment path between two sequences of feature vectors $X \in \mathbb{R}^{M \times D}$ and $Y \in \mathbb{R}^{N \times D}$, in which M and N are the lengths of vectors X and Y respectively and D is the dimension of the feature vector. For example: for chroma vectors, $D = 12$. An alignment path is defined as two nondecreasing sequences $p \in \mathbb{N}^L$ and $q \in \mathbb{N}^L$, that satisfy the following three conditions:

- Boundary condition: $p[1] = 1, q[1] = 1, p[L] = N$ and $q[L] = M$;
- Monotonicity condition: $p[1] \leq p[2] \leq \dots \leq p[L]$ and $q[1] \leq q[2] \leq \dots \leq q[L]$;
- Step size condition: $(p[l + 1] - p[l], q[l + 1] - q[l]) \in \{(1, 1), (1, 0), (0, 1)\}$ for $l \in [1, L - 1]$.

In the alignment path p, q , $p[i] = n$ and $q[i] = m$ implies that the n^{th} feature vector in X is matched to the m^{th} feature vector in Y . The optimal alignment path is the path with the lowest total cost, in which the total cost is defined as the sum of local costs c between each pair of features $(X[p[i]], Y[q[i]])$ on the alignment path. Finding the optimal path p, q can thus be defined as the following minimization problem:

$$p, q = \underset{p, q}{\operatorname{argmin}} \sum_{i=1}^L c(X[p[i]], Y[q[i]]) \quad (5.1)$$

As an example, consider two sequences $X = [0, 1, 2, 3, 2, 1]$ and $Y = [1, 2, 3, 2, 0]$ and define the cost function as the absolute distance, i.e. $c(X[p[i]], Y[q[i]]) = |X[p[i]] - Y[q[i]]|$. The optimal alignment path is illustrated in Figure 5.2. In this path, the first two elements of X are aligned to the first element of Y ; the third element of X is aligned to the second element of Y ; the fourth element of X is aligned to the third element of Y , etcetera. Therefore, the optimal path is $p, q = [1, 2, 3, 4, 5, 6], [1, 1, 2, 3, 4, 5]$. This path has length $\sum_{i=1}^6 c(X[p[i]], Y[q[i]]) = c(0, 0) + c(1, 0) + c(2, 2) + c(3, 3) + c(2, 2) + c(1, 0) = 0 + 1 + 0 + 0 + 0 + 1 = 2$.




Figure 5.2: Alignment of $X = [0, 1, 2, 3, 2, 1]$ and $Y = [1, 2, 3, 2, 0]$

To determine the optimal alignment path, one could compute the total cost of all possible warping paths and then return the path with the lowest cost. Unfortunately this is not feasible for large sequences, because the number of possible warping paths is exponential in the length of the sequences. Luckily, there exist other methods to find the optimal alignment path.

Consider a path p_n, q_m with length k that is a prefix of the alignment path. That is: p_n, q_m fulfills all requirements of the alignment path, except for the last part of the boundary criterion. So:

- $p_n[1] = 1, q_m[1] = 1, p_n[k] = \mathbf{n} \in [1, N]$ and $q_m[k] = \mathbf{m} \in [1, M]$;
- $p_n[1] \leq p_n[2] \leq \dots \leq p_n[k]$ and $q_m[1] \leq q_m[2] \leq \dots \leq q_m[k]$;

- $(p_n[l+1] - p_n[l], q_m[l+1] - q_m[l]) \in \{(1, 1), (1, 0), (0, 1)\}$ for $l \in [1, k-1]$.

Now note that, thanks to the step size condition, any path p_n, q_m starts with either p_{n-1}, q_{m-1} ; p_n, q_{m-1} or p_{n-1}, q_m , followed by one additional step. Therefore, if we know the costs of the optimal p_{n-1}, q_{m-1} ; p_n, q_{m-1} and p_{n-1}, q_m , we can derive the cost of the optimal path p_n, q_m in constant time, using the following equation:

$$\begin{aligned} \text{cost}(p_n, q_m) &= \min_{p_n, q_m} \sum_{i=1}^k c(X[p_n[i]], Y[q_m[i]]) \\ &= c(X[p_n[l]], Y[q_m[l]]) + \min \begin{cases} \sum_{i=1}^{k-1} c(X[p_{n-1}[i]], Y[q_{m-1}[i]]) \\ \sum_{i=1}^{k-1} c(X[p_n[i]], Y[q_{m-1}[i]]) \\ \sum_{i=1}^{k-1} c(X[p_{n-1}[i]], Y[q_m[i]]) \end{cases} \end{aligned} \quad (5.2)$$

All paths p_n, q_1 have the special property that they must exclusively consist of horizontal steps, i.e. $(p_n[l+1] - p_n[l], q_1[l+1] - q_1[l]) = (1, 0)$ for $l \in [1, k-1]$. Therefore, we can calculate their cost without minimization:

$$\begin{aligned} \text{cost}(p_n, q_1) &= \min_{p_n, q_1} \sum_{i=1}^k c(X[p_n[i]], Y[q_1[i]]) \\ &= \sum_{i=1}^n c(X[p[i]], Y[q[1]]) \end{aligned} \quad (5.3)$$

Likewise, we do need to minimize for calculating the cost of any path p_1, q_m :

$$\begin{aligned} \text{cost}(p_1, q_m) &= \min_{p_1, q_m} \sum_{i=1}^k c(X[p_1[i]], Y[q_m[i]]) \\ &= \sum_{i=1}^m c(X[p[1]], Y[q[i]]) \end{aligned} \quad (5.4)$$

Also, note that the optimal path p_n, q_m with $n = N$ and $m = M$ equals the optimal full alignment path p, q .

The Dynamic Time Warping algorithm uses Equations 5.2, 5.3 and 5.4 to efficiently compute the optimal full alignments path p, q and its cost. Basically, the algorithm computes the cost of any path p_n, q_m with $n \in [1, N]$ and $m \in [1, M]$ in a convenient order and stores the result. This way, each of these $(N \times M)$ costs is computed exactly once, so the algorithm runs in $\mathcal{O}(N \times M)$ time.

The pseudocode of the DTW algorithm is given in Algorithm 1. The algorithm takes a cost matrix C with $C[n, m] = c(X[n], Y[m])$. Any entry $D[n, m]$ contains the cost of the optimal path p_n, q_m , so after running the algorithm $D[N, M]$ contains the minimum cost of p, q . We can find the optimal alignment path p, q by following the arrows backwards from $P[N, M]$.

Algorithm 1 Pseudocode of the DTW algorithm

```

1: function DTW(Cost Matrix  $C$  of size  $N \times M$ )
2:   Initialize Accumulated cost Matrix  $D$  of size  $N \times M$ 
3:   Initialize Path Matrix  $P$  of size  $N \times M$ .
4:    $D[1, 1] = C[1, 1]$ 
5:    $P[1, 1] = \cdot$ 
6:   for all  $n \in [2, N]$  do
7:      $D[n, 1] = C[n, 1] + D[n - 1, 1]$ 
8:      $P[n, 1] = \leftarrow$ 
9:   end for
10:  for all  $m \in [2, M]$  do
11:     $D[1, m] = C[1, m] + D[1, m - 1]$ 
12:     $P[1, m] = \downarrow$ 
13:  end for
14:  for all  $n \in [2, N]$  do
15:    for all  $m \in [2, M]$  do
16:       $s = \min D[n - 1, m - 1], D[n, m - 1], D[n - 1, m]$ 
17:      if  $s \equiv D[n - 1, m - 1]$  then
18:         $D[n, m] = C[n, m] + D[n - 1, m - 1]$ 
19:         $P[n, m] = \swarrow$ 
20:      else if  $s \equiv D[n, m - 1]$  then
21:         $D[n, m] = C[n, m] + D[n, m - 1]$ 
22:         $P[n, m] = \downarrow$ 
23:      else
24:         $D[n, m] = C[n, m] + D[n - 1, m]$ 
25:         $P[n, m] = \leftarrow$ 
26:      end if
27:    end for
28:  end for
29: end function

```

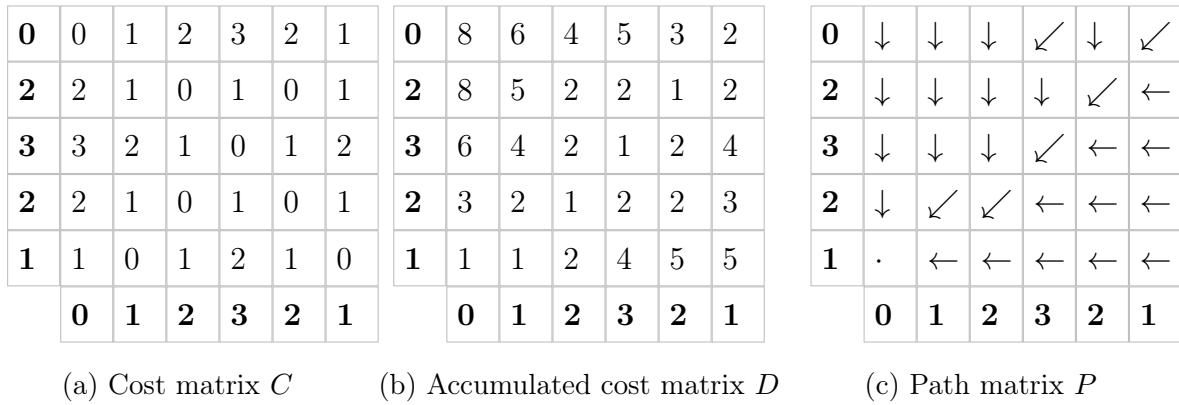


Figure 5.3: Cost matrix and DTW output for the alignment of our toy example $X = [0, 1, 2, 3, 2, 1]$ and $Y = [1, 2, 3, 2, 0]$

Let us now test the DTW algorithm on our example $X = [0, 1, 2, 3, 2, 1]$ and $Y = [1, 2, 3, 2, 0]$ from Figure 5.2. First, we calculate the cost matrix C . Remember that in our example, c is defined as the absolute difference. For example, $C[4, 1] = c(X[4], Y[1]) = |X[4] - Y[1]| = |3 - 1| = 2$. The resulting cost matrix C is given in Figure 5.3a. Figures 5.3b and c show the accumulated cost matrix D and the path matrix P respectively, which are calculated by the algorithm. We indeed find the optimal alignment path $(p, q = [1, 2, 3, 4, 5, 6], [1, 1, 2, 3, 4, 5])$ by following the arrows back from $P[N, M]$ and read from $D[N, M]$ that its total cost truly equals 2.

Dynamic Time Warping in music alignment

In the previous subsection, we saw an example of DTW on a one-dimensional signal. When aligning two music representations, the **features** typically have more dimensions. For example, chroma features would be a sensible choice: they can be calculated from both the audio and the synthesized MIDI. In that case, the features are twelve-dimensional. Chroma features are for example used for alignment by Hu et al. (2003). Prätzlich and Müller (2016) and Wang et al. (2014) also use chroma based-features, but those are combined with features capturing onset information. Other methods, those designed by Dixon and Widmer (2005) and Raffel and Ellis (2015), use the result from a Constant-Q transform.

As we have seen in Section 2.5.1, we need to set a **time scale** for the feature vectors. Some methods, for example Dixon and Widmer (2005), compute feature vectors over short, overlapping frames of audio. Other methods, like Raffel and Ellis (2015) use beat-synchronous feature vectors, obtained by aggregating the original feature vectors between two beats.

Multidimensional features require an appropriate **cost function**. A common choice (used by e.g. Turetsky and Ellis (2003); Wang et al. (2014); Raffel and Ellis (2015)) is the cosine

distance. The cosine distance between two d -dimensional features x and y is defined as:

$$c(x, y) = 1 - \frac{\langle x|y \rangle}{\|x\| \cdot \|y\|} = 1 - \frac{\sum_{i=1}^d x_i \cdot y_i}{\sqrt{\sum_{i=1}^d x_i^2} \cdot \sqrt{\sum_{i=1}^d y_i^2}} \quad (5.5)$$

Some methods use a **penalty** to discourage non-diagonal moves (Raffel and Ellis, 2015). That is, Equation 5.1 is replaced by Equation 5.6, in which $\Phi(i) \leq 0$ if $p[i] = p[i - 1]$ or $q[i] = q[i - 1]$.

$$p, q = \underset{p, q}{\operatorname{argmin}} \sum_{i=1}^L c(X[p[i]], Y[q[i]]) + \Phi(i) \quad (5.6)$$

In music alignment, it is not always necessary that the alignment path p, q spans the entirety of feature vectors X and Y . For instance, sometimes we want to match only a part of a MIDI file to the audio recording. In order to enable subsequence matching, we drop the boundary condition ($p[1] = 1, q[1] = 1, p[L] = N$ and $q[L] = M$) and relax it to the more flexible condition that either $g \cdot N \leq p[L] \leq N$ or $g \cdot M \leq q[L] \leq M$. g is the **gully** parameter and determines the proportion of the subsequence that must be successfully matched. Raffel and Ellis (2015) for example choose a gully of 0.95, allowing for some tolerance that the beginning or ending of a MIDI file is incorrect.

To conclude, some methods place **global constraints** (Müller, 2015, Chapter 3) on the alignment path. These global constraints reduce complexity of the DTW algorithm and prevent paths that diverge too much from the diagonal line. A global constraint only allows points on the warping path inside the global constraint region $R \subseteq [1 : N] \times [1 : M]$. Two well-known constraints are the Sakoe-Chiba band (Sakoe and Chiba, 1978) and the Itakura parallelogram (Itakura, 1975). The Sakoe-Chiba band constrains the path to lie within a fixed distance of the diagonal, as shown in Figure 5.4a. The Itakura parallelogram is illustrated in Figure 5.4b and constrains the path to lie within a parallelogram around the diagonal of the matrix.




Figure 5.4: Global constraint regions

5.1.2 Selected system

As we have seen in Section 5.1.1, Dynamic Time Warping (DTW) is a common technique to align two feature vectors, for example two representations of the same song. We

Parameter	Setting
Feature representation	log-magnitude Constant-Q transform
Time scale	every 46 milliseconds
Cost function	cosine distance
Penalty	median distance of all pairs of frames
Gully	0.96
Band path constraint	none

Table 5.1: Parameter settings of audio-MIDI DTW alignment algorithm

also learned that there are many variations of the DTW algorithm: one can use different features, time scales, cost functions, penalties, gully parameters and global constraints. In this section, I will explain the MIDI-to-audio alignment method I selected for DECIBEL.

For alignment between MIDI files and audio recordings, DECIBEL uses a DTW algorithm by Raffel and Ellis (2016b). I selected this algorithm for a couple of reasons. In the first place, it benefits from the two advantages of DTW: the algorithm calculates an easily interpretable alignment confidence score, which gives a good indication of the alignment quality. Furthermore, the algorithm is conceptually simple and easy to implement. Typically, the performance of DTW heavily relies on the chosen parameters. An advantage of this specific system is that the optimal parameter setting is trained on a synthetically trained data set of 1000 MIDI files which were transcriptions of Western popular music songs, using Bayesian optimization.

Before we look at the specific parameter settings, let us first consider the **outline** of the algorithm. The outline is illustrated with an example in Figure 5.5. First, all MIDI files are synthesized using the fluidsynth¹ software synthesizer with the FluidR3_GM soundfont. Now we have a waveform representation for both the audio (Figure 5.5a) and the MIDI file (Figure 5.5b). Note that our example MIDI file starts with silence, while in the audio recording the music starts immediately. Also, the MIDI file has a longer duration, as the MIDI file repeats the chorus an additional time, compared to the audio file. Then, the Constant-Q transform is calculated for both the audio (Figure 5.5c) and the synthesized MIDI waveform (Figure 5.5d). Features are found by aggregation over the Constant-Q transform vectors. Then, the optimal path between the audio file and the synthesized MIDI is calculated using DTW. This results in an optimal path and the alignment confidence score, as illustrated in Figure 5.5e. In this figure, we see that the alignment path starts not in the coordinate (0, 0), but a bit to the right: the silence at the start of the MIDI file is not mapped to any position in the audio file. The same goes for the end of the MIDI file, which is a superfluous repetition of the chorus. Finally, this alignment path is used to remap the MIDI file to the audio recording (see Figure 5.5f).

DECIBEL uses the unchanged **parameter setting** reported in the paper by Raffel and Ellis (2016b). The parameters are listed in Table 5.1. I did not experiment with alternative parameter settings as the parameters were already optimized for a dataset in the same genre as mine. In the optimal parameter setting, the features are represented by log-magnitude Constant-Q transform. This is the Constant-Q transform we have seen in

¹<http://www.fluidsynth.org/>

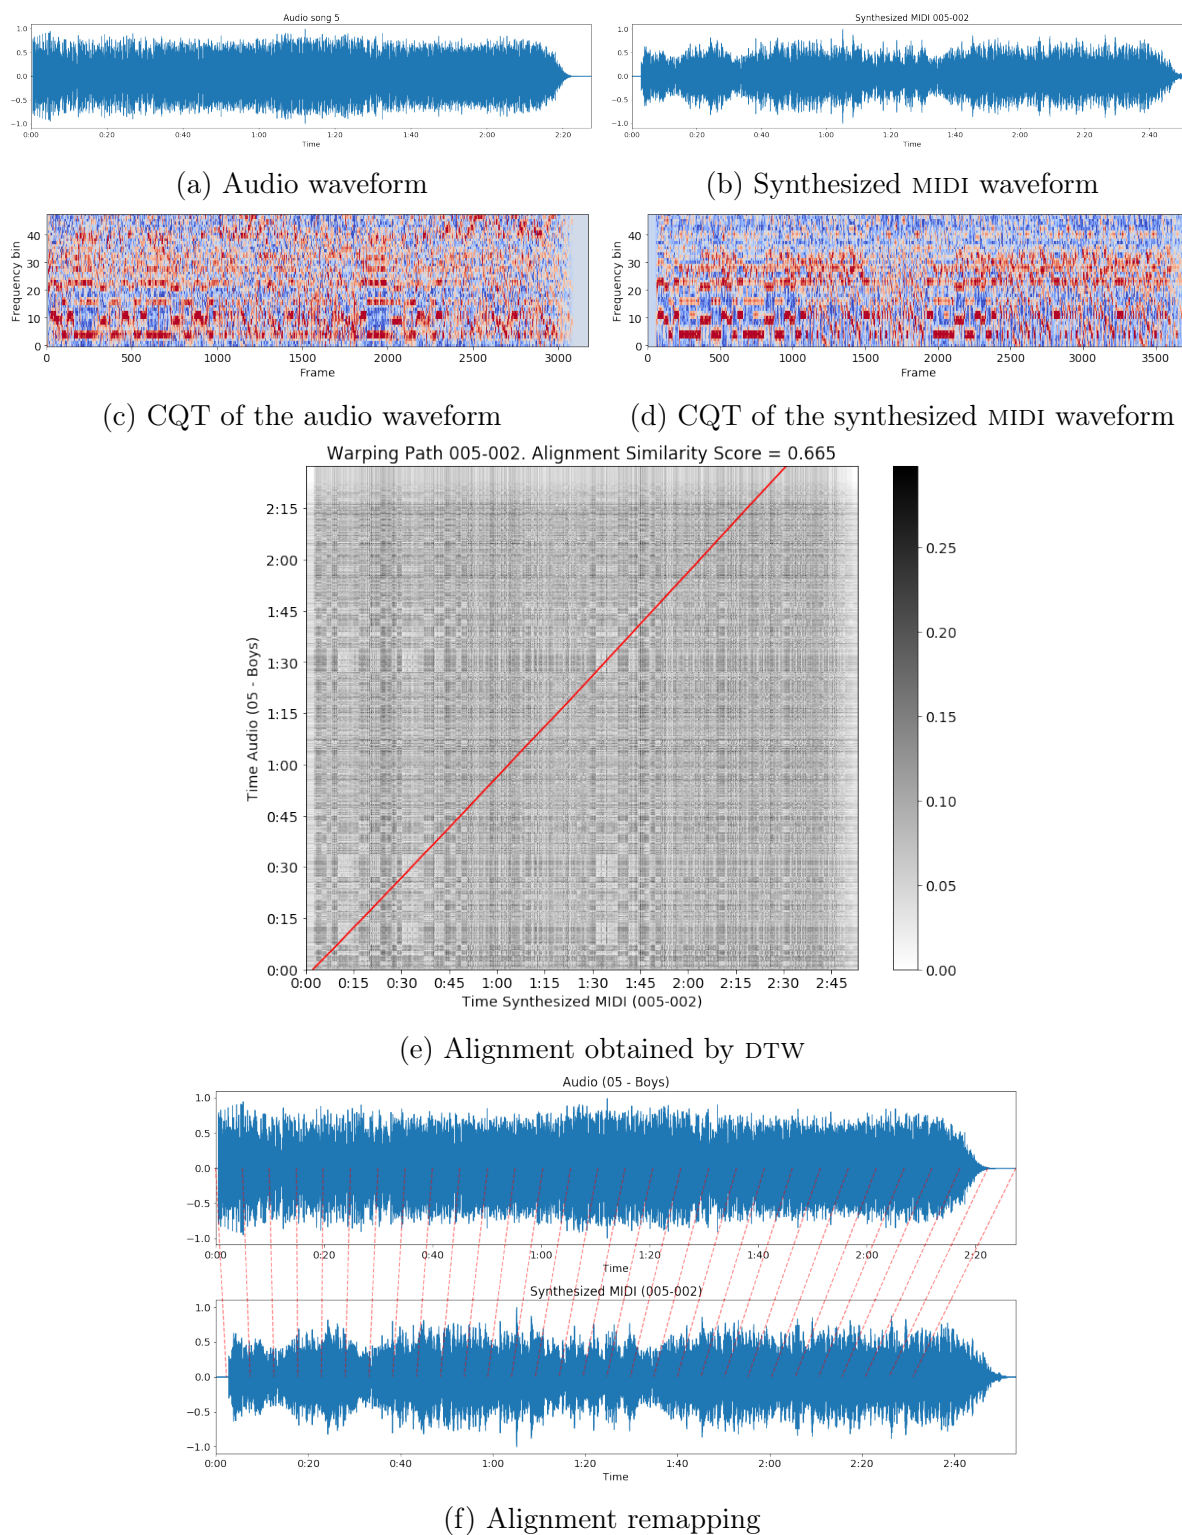


Figure 5.5: Illustration of the alignment procedure

Section 2.5.1, but the log of the features is calculated to mimic human perception more closely. An optimal hop size of 1024 samples is found, which corresponds to a time scale of 46 milliseconds, given the sampling rate of 22050 Hz: $\frac{1024}{22040} = 0.046\text{s}$. Feature vectors are normalized by L2 norm before calculating the local distances. This is equivalent to using the cosine distance. A penalty of the median distance between all pairs of frames is found to give the best results. In the optimal system, the gully parameter is 0.96 and there is no band path constraint.

The **output** of the DTW system is an optimal path and its alignment confidence score. The path specifies which time point, measured in seconds, in the MIDI file is aligned to which time point in the audio file. We use this path to recompute the times in the MIDI file using the `pretty_midi` package (Raffel and Ellis, 2014). This gives us the “audio-aligned MIDI file” we will use in the chord estimation step. The alignment confidence score is the mean distance between all pairs of frames over the entire aligned portion of both feature sequences. A qualitative evaluation on 500 real word MIDI/audio pairs by Raffel and Ellis (2016b) shows that the alignment confidence score is a reliable measure for the quality of the alignment in most cases, although there were a few outliers which had a low (good) alignment confidence score despite being aligned badly. In general, MIDI/audio pairs with an alignment confidence score below 0.85 are aligned well. In the next section, we will do a similar evaluation for to verify if the alignment confidence score is a good indicator for the alignment quality in our data set.

5.1.3 Evaluation of MIDI-to-audio alignment

After aligning all MIDI files in the data set to the corresponding audio file, I found an average alignment confidence score of 0.768. In order to verify if the results found by (Raffel and Ellis, 2016a) are applicable to my data set, I evaluated the performance of the DTW system on a random sample of 25 MIDIS. For each MIDI file, I synthesized the realigned MIDI version and played it simultaneously with the original audio file in Sonic Visualiser, listening to the realigned MIDI file on the left earphone and the original audio on the right earphone. In this listening test, I classified each MIDI into one of three alignment quality categories: bad alignments; alignments with minor issues; or good alignments. The results of this evaluation can be found in the Appendix, in Table B.1 and is also depicted in a violin plot in Figure 5.6. A violin plot (Hintze and Nelson, 1998) is an alternative for the box plot that reveals density information from the data. My evaluation confirmed Raffel’s observation that alignments with a low alignment confidence score are good, while high alignment score (above 85%) have major issues. These problems are mostly due to MIDI files in a wrong transposition or MIDI files that were bad transcriptions, for example because they only represented one part of the song.

5.2 Chord estimation on MIDI

In the previous section we saw how DECIBEL aligns MIDI files to audio recordings. As a next step, we need a chord estimator to calculate the chord sequences from this realigned




Figure 5.6: Violin plot showing the distribution of alignment confidence scores for each rating in our qualitative evaluation.

MIDI file. In this section, I give an overview of related work on MIDI ACE in Subsection 5.2.1 and describe DECIBEL’s MIDI chord estimation method in Subsection 5.2.2. This method will be briefly evaluated in Subsection 5.2.3.

5.2.1 Related work on chord estimation on MIDI

Chord estimation in MIDI files² is the task of dividing a MIDI file into segments, in such a way that each segment boundary corresponds to a chord change, and assigning a chord label to each segment. In contrast to audio chord estimation, only few methods have been proposed that extract chords from symbolic music like MIDI.

In a pioneering paper, Winograd (1968) proposes a grammar-based approach to perform automatic tonal (roman numeral) analysis. This type of analysis identifies chords in their harmonic context. Maxwell (1992) developed a rule-based expert system that performs harmonic chord function analysis for tonal music, for example Bach keyboard pieces. Temperley and Sleator (1999) present a computational system for analyzing both metrical and harmonic structure. The system is based on eight preference rules, which are criteria for selecting the best analysis of a piece out of many possible ones. Their algorithm is implemented in the first version of Melisma Music Analyzer (Sleator and Temperley, 2001).

Pardo and Birmingham (2001, 2002) propose another method for segmentation and chord labeling. They segment the score using partition points (the set of all start and end points, excluding duplicates). Then, the score for each segment is calculated by computing the minimum of the distance between the segment and any of the 180 (template, root) pairs. An example of a template is $\langle 0, 3, 7 \rangle$, corresponding to a major chord. The authors give 15 different templates and 12 possible roots - one for each pitch class. In case of a tie, the best (template, root) pair is found using tie breaking rules. Inspired by aforementioned

²also referred to as *symbolic chord recognition* or MIDI ACE

research by Pardo and Birmingham, Scholz and Ramalho (2008) propose COCHONUT: a new way to recognize chords from symbolic MIDI guitar data in bossa nova music, which is also dealing with complex chords like ninth and suspended chords. Chord estimation is split into three steps: (1) run a segmentation algorithm; (2) apply a utility function to identify the most probable chords for each segment and build a graph representing them; and (3) choose the best chord label for each segment, considering contextual information.

Raphael and Stoddard (2003, 2004) use probabilistic models to perform functional harmonic analysis on MIDI data. The analysis is performed on a fixed musical period, for example a measure. They use 12 pitch classes. A chord is specified by the tonic, mode and chord. For example, $(t, m, c) = (2, \text{major}, \text{II})$ would represent a triad in the key of D major built on the second scale degree. So this chord contains the pitches E, G and B (e minor). The most likely chord sequence is computed with a HMM using the Baum-Welch algorithm.

Radicioni and Esposito (2010) perform chord estimation using a HMPerceptron model, in which the domain knowledge is modeled in Boolean features. As chord vocabulary, they consider chords with 12 possible root nodes, 3 possible modes and 3 possible added notes, resulting in 108 possible chord labels.

In more recent research, Masada and Bunesu (2017) propose a machine learning model for chord estimation that uses semi-Markov Conditional Random Fields (semi-CRFs) to perform a joint segmentation and labeling of symbolic music.

In this section, we have seen that the number of algorithms for MIDI chord recognition is limited. The existing algorithms are not directly usable in DECIBEL, for the following reasons:

- The algorithm is designed for a specific music genre (usually classical music);
- The algorithm is based on machine learning. Training the system requires a lot of labeled data, which is not available for MIDI files of popular music; or
- The algorithm recognizes functional harmony instead of chord labels.

5.2.2 Implementation of CASSETTE

In Section 5.2.1 we have seen that the existing algorithms for MIDI chord recognition are not suitable for direct use in our DECIBEL system. Therefore, I designed CASSETTE (Chord estimation Applied to Symbolic music by Segmentation, Extraction and Tie-breaking TEmplate matching). CASSETTE is a template-matching based algorithm for MIDI chord recognition that is easy to implement and understand. Similar to the good old cassette tapes, this algorithm is certainly not state of the art. However, it is simple to implement and does not require any training.

CASSETTE recognizes chords in a three-step procedure: first, it segments each audio-aligned MIDI file. Then, it calculates a weighted chroma feature for each of the segments,

based on the notes that are present within the segment. Finally, CASSETTE matches the features of each segment to the features of a predefined chord vocabulary and assigns each segment to the most similar chord. In the remainder of this section, we zoom in on each of these three steps.

CASSETTE **segments** the MIDI both on the bar and on the beat level. In many cases, segmentation on the bar level is sufficient, as chord changes in popular music often are placed on the downbeat, i.e. on the first beat of a bar. An advantage on segmentation on the bar level is that non-harmony notes, which are typically short, are less problematic in the template matching step as they have relatively lower weights than the (typically longer) harmony notes. On the other hand, segmentation on the bar level does not work well for songs that have chord changes at other positions than the start of a bar. That is why CASSETTE segments the MIDI file on the beat level as well. This typically results in a chord sequence with more chord changes, although some of which are based on non-harmony notes. For segmentation, CASSETTE uses the `pretty_midi` package by Raffel and Ellis (2014). Bar segments are found with the `get_downbeats` function, while the algorithm finds beat segments with the `get_beats` function.

After segmentation, the next step is **feature extraction** on each of the segments of the MIDI file. For each segment, CASSETTE extracts the notes sounding between its start and end time. From these notes, we calculate weighted chroma, a feature that is similar to A-weighted chroma as used by Bonvivi (2014, Section 4.1). In order to calculate the weighted chroma feature for a given segment, CASSETTE computes for each note the product of the MIDI velocity and the proportion of the bar during which this note sounds, and sums this product over all notes in the same pitch class. For example, the weighted chroma of a quarter note C in a 4/4 bar with MIDI velocity of 100 is $[25, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]$. In the resulting vector, both the louder notes (with high velocity) and the longer notes (with higher duration) in the MIDI file are relatively more important, compared to softer or shorter notes. CASSETTE normalizes the weighted chroma vector by dividing each element by the total sum of all its elements. This makes the feature invariant to the total loudness and duration of the notes in the segment.

As a third and final step, we need to find the **best matching chords** for each segment. Therefore, CASSETTE calculates the template similarity score between the feature of the segment and the feature of each template in our chord vocabulary. In this project, I use a vocabulary of 24 chords, consisting of all 12 major chords and all 12 minor chords, and the no-chord. The chroma-like feature of each template is a 12-dimensional vector, in which each value is 1 if the corresponding note occurs in that chord and the value is 0 otherwise. For example: the D minor chord consists of a D, F and an A. The corresponding chroma would be $[0, 0, 1, 0, 0, 1, 0, 0, 0, 1, 0, 0]$. The score function that measures the similarity between the chroma of the segment and the chroma template is based on earlier work by Pardo and Birmingham (2002). The template similarity score S is calculated with the formula $S = P - (N + M)$. P is the positive evidence: the sum of the weights of the chroma of the bar which match a template element. N is the negative evidence: the sum of the weights of the chroma of the bar which does not match a template element. M stands for misses: the count of template elements not matched by any note. High scores correspond to well-matched templates. For each bar, CASSETTE finds the chord with the highest template similarity score. If the score is -3 or lower, the algorithm assigns a no-

Sementation method	wCSR	OvS	UnS	Seg
Beat	0,793	0,832	0,881	0,805
Bar	0,700	0,953	0,658	0,658

Table 5.2: Results of MIDI chord recognition for the 50 MIDI files with the lowest alignment error.

chord, as there is no evidence for any chord if three notes or more from the template are missing. Furthermore, CASSETTE applies a single tie-breaking rule: if multiple templates have the same template similarity score, it selects the template whose root pitch has the greatest weight in the segment’s chroma.

5.2.3 Evaluation of CASSETTE

In the previous section, we have seen how CASSETTE estimates chords from MIDI files. For the evaluation of the system, we run into a minor issue: a suitable data set with chord annotations for MIDI files of popular music does not yet exist. As a solution, I selected 50 well-aligned MIDI files from my data set and tested them against the Isophonics annotations, as described below.

First, I needed a subset of MIDI files for which the timing information was consistent with the Isophonic annotations. In Section 5.1.3 we observed that MIDI/audio pairs that are aligned well typically have a low alignment error. Therefore, I selected the 50 MIDI files with the lowest alignment error. This subset had alignment confidence scores ranging from 0.459 to 0.613, so we can reasonably expect that these MIDI files are well aligned to the audio.

Then, I ran the CASSETTE algorithm on the audio-aligned versions of these 50 MIDI files and calculated the wCSR for each of the resulting chord label sequences. The results are shown in Table 5.2. Note that particularly the beat-based MIDI chord recognition method performs quite well in terms of wCSR with a 79.3% score. Compared to the bar-based MIDI chord recognition method, beat-based chord recognition has minor oversegmentation issues, while the bar-based method is undersegmenting. This also explains the lower wCSR score of 70.0% of the bar-based method. We can conclude that CASSETTE, despite its simplicity, performs quite well on our data set with popular music and a limited chord vocabulary.

Another interesting property of CASSETTE is that it computes template similarity scores, which we can use to calculate an approximation score for the quality of the output chord sequences. The template similarity score for each segment gives an indication of the fit of the chord label to this segment: a high score indicates a well-fitting chord label while a low-score segment probably has a less suitable chord label. This may be the case when the actual chord label is more complex than the best fitting chord in our simplified vocabulary. When we average the scores over all segments, we get an approximation score for the quality of the chord label. We call this score the Average Template Similarity (ATS) and we will see its advantages in the next section.

5.3 MIDI file selection

In the previous section we saw that CASSETTE, our MIDI chord recognition method, was quite successful for a 50-song subset of well-aligned MIDI files. However, when I repeated the aforementioned experiment on the full data set, results were considerably worse: when comparing the chord sequences found by MIDI file alignment and chord recognition system on all MIDI files, I found a WCSR of 59.9% for bar-based chord recognition and 62.6% for beat-based chord recognition. These results are certainly not competitive to the performance of other (audio) chord recognition systems. However, we have already seen in Section 5.1 that some MIDI files are badly aligned. It is not very surprising that these files will not yield good chord labels. Therefore, I discarded all MIDI files with an alignment error higher than 0.85. As expected, we see a great shift in performance: the WCSR on all MIDI files that are sufficiently aligned, is 69.2% (bar) and 72.5% (beat).

After this MIDI selection step based on alignment quality, we have 565 MIDI files left. There are still multiple MIDI files per song. We can do an additional selection step to find the best MIDI file of the song. For each song, the MIDI file that is truly the best, is the one that yields the highest CSR. However, we cannot calculate the CSR for unlabeled data. Therefore, we need a proxy measure.




Figure 5.7: Scatter plot comparing CSR and ATS. We see that there exists a correlation between CSR and ATS for both beat and bar segmentation. Therefore, ATS may be a good proxy measure for CSR.

As shown in Figure 5.7, there is a correlation between the CSR and the ATS (Average Template Similarity) that is calculated by CASSETTE. The Pearson correlation coefficient is 0.481 for beat segmentation and 0.598 for bar segmentation. Figure 5.8 shows the estimated ranking of MIDI files based on ATS, compared to the real ranking based on CSR. A ranking of 1 corresponds to the (estimated) best MIDI file, a MIDI file with rank 2 is the second best, etcetera. The red numbers in each of the cells represent the number of songs with this estimated ranking based on ATS (x -axis) and CSR (y -axis). For example:

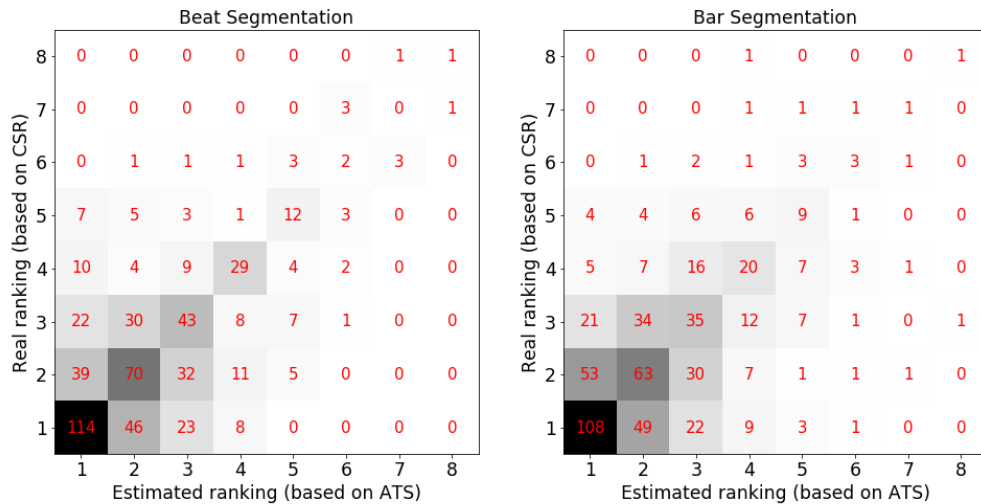


Figure 5.8: Estimated ranking (based on ATS) compared to real ranking (based on CSR) for both beat and bar segmentation. The red numbers in each of the cells correspond to the number of songs with this estimated ranking based on ATS (x -axis) and CSR (y -axis).

46 songs are ranked first based on CSR, while they are ranked second based on ATS. As we see in the figure, in many songs the MIDI file with the highest CSR equals the MIDI file with the highest ATS. For most other songs, the CSR of the MIDI with the highest ATS is just a bit worse than the actual best MIDI, as shown in Figure 5.9. In this scatter plot, each point corresponds to a song from our data set. We see the CSR of the estimated best MIDI of each song on the x -axis and the CSR of the truly best MIDI of the same song on the y -axis. Points on the red line (where $x = y$) correspond to songs for which the best MIDI file was estimated correctly. As we can see, there are very few songs for which there is a big difference between CSR of the best MIDI and CSR of the estimated best MIDI. There are three songs for which the difference is higher than 0.5. In two of these songs, the estimated best MIDI was a semitone transposed, compared to the original audio. In the third song, the estimated best MIDI is only a transcription of part of the song. However, in the vast majority of songs, the difference between the estimated and truly best MIDI files are quite small: most points are close to the red line. We can therefore conclude that the ATS is a suitable score measure to select the estimated best MIDI file for each song.

Finally, the performance of our selection method on all songs in the data set is summarized in Table 5.3. If we would not select the estimated best MIDI file, but just calculate the WCSR over all MIDI files for all songs, the resulting WCSR would be 72.5% for beat segmentation and 69.2% for bar segmentation. This is the baseline for our selection method. On the other hand, if we would know beforehand what would be the MIDI with the best CSR and select this MIDI file, the WCSR of chord sequences from exclusively the best MIDIs would be 79.5% and 75.5%. This is an upper boundary for the selection method. We now see that the chord sequences from MIDIs selected with our proposed selection method (Best-ATS) have WCSRs of 75.7% and 72.9% and thereby beat the baseline with more than three percentage points. In addition, each of the three segmentation measures is improved by selecting the MIDI with the highest ATS.




Figure 5.9: CSR of the real best MIDI compared to the CSR of the estimated best MIDI for both beat and bar segmentation. Points on the red line (i.e. $x = y$) correspond to songs for which the best MIDI file was estimated correctly. The distance in y direction between each point and the red line is the difference between the CSR of the best MIDI and the CSR of the estimated best MIDI file.

	WCSR		OvSeg		UnSeg		Seg	
	Beat	Bar	Beat	Bar	Beat	Bar	Beat	Bar
All MIDIs	0.725	0.692	0.803	0.931	0.866	0.709	0.769	0.703
Best-ATS MIDI	0.757	0.729	0.813	0.932	0.872	0.733	0.776	0.725
Best-CSR MIDI	0.795	0.755	0.826	0.937	0.880	0.734	0.794	0.725

Table 5.3: Performance comparison of three MIDI selection methods: no selection, MIDI file with best ATS, MIDI file with best CSR.

5.4 Conclusion

In this section, we have examined the MIDI subsystem of DECIBEL. This subsystem extracts chord sequences from MIDI files by first aligning the MIDI file to the audio file and then running a simple chord estimation method on the re-aligned MIDI file. In our data set, we have collected multiple MIDI files for each song. The subsystem computes a chord sequence for each MIDI file, so we obtain multiple chord sequences per song. We have seen that we can select “good” (and in many cases the best) MIDIs for each song by (1) ignoring MIDIs with a high alignment error and (2) selection of the MIDI with the highest Average Template Similarity (ATS). This resulted in WCSRs of 75.7% for the beat segmentation and a 72.9% for bar segmentation.

Chapter 6

Automatic Chord Estimation on tabs

In the previous chapters we have seen how our system estimates audio-timed chord labels from audio representations and MIDI files. In this chapter, we will look at DECIBELs third subsystem, which uses tab files for estimating chord labels. As we have seen in Section 2.5.3, there exist two types of tabs and we can easily extract chord symbols from both of them: guitar tablature indicates the instrumental fingering, which directly maps to the notes of a chord; in a chord sheet, the chord symbols are represented explicitly, together with the lyric syllables with which they have to be timed. In this chapter, I will explain how DECIBEL uses tabs from either of these two types in order to estimate audio-timed chord labels. The outline of this subsystem is also illustrated in Figure 6.1. First, the tabs are parsed and the chord information is extracted. This step is described in Section 6.1. As a next step, DECIBEL aligns the chord information to the audio file, as described in Section 6.2. This gives us multiple chord estimates: one for each tab file of the corresponding song. In Section 6.3, I explain how DECIBEL selects the expected best tab file for each song. Finally, I summarize the chapter in Section 6.4.

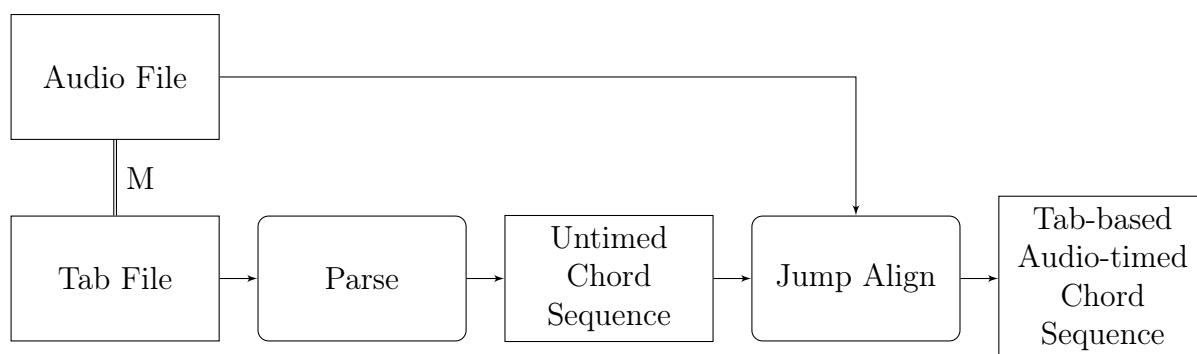


Figure 6.1: Diagram of DECIBEL’s tab subsystem; the M between Audio File and Tab File indicates that they are manually matched.

6.1 Tab parsing

Before DECIBEL can align the tab files to the audio, it first needs to parse them and extract the chord information. The used tab parsing method is explained in Section 6.1.1 and evaluated in Section 6.1.2.

6.1.1 Tab parsing methodology

DECIBEL parses the tabs in a similar way to the parser proposed by Macrae (2012). First, it classifies each line in the tab file to a line type. Then, it segments the tab by splitting on empty lines. As a next step, all systems in each segment are identified. We define a system as a set of subsequent lines that belong to each other. For example: a tab system is very common in guitar tablature files and consists of exactly six subsequent tab files. In chord sheets, a common system is the alteration between chord lines and lyrics lines. From these systems, DECIBEL can then extract the chord labels. Thereby, the system retains line information (i.e. the line of the chord in the text file), as this will be used in the tab-audio alignment step. The steps of line type classification, segmentation and system and chord extraction are explained more thoroughly in the remainder of this section.

Line type classification

For each line in the tab file, DECIBEL estimates its line type using a set of heuristics. We distinguish the following line types: *empty*, *chords*, *tuning definition*, *capo change*, *structural marker*, *chord definition*, *tablature*, *lyrics* and *combined chord and lyrics*.

- A line has the *empty* line type if it is empty or consists only of a space.
- To determine if a line has the *chords* line type, the system first splits the line by spaces. Then it checks if each element of the line matches the chord pattern. An element is a chord if:
 - It consists of at most 10 characters;
 - The first character is a note letter (a, b, c, d, e, f or g);
 - The element does not contain sequences of four digits (as this would indicate a chord definition);
 - The element does not contain symbols; and
 - Any three letter sequence is in the following list: *min*, *add*, *aug*, *dim*, *maj*, *sus*, *flat*.
- A line has the *tuning* line type if it contains the word “tuning”.
- To find out if a line is a *structural marker*, our parser searches for words like “verse”, “chorus” or “bridge”.

- A line is a *chord definition* if it contains a sequence of exactly 6 digits.
- A line is a *tablature line* if it contains at least 10 characters which are a digit, hyphen, vertical bar, slash, letter ‘h’, ‘b’ or ‘p’ or a space and the number of hyphens is larger than the number of spaces.
- A *lyrics line* fulfills each of the following conditions:
 - It does not contain square brackets, the equals sign or the at sign;
 - It contains at most 10 hyphens;
 - Either:
 - * It consists of just one word, which has only letter characters and contains at least three of the same letters after each other, like ‘ooooh’ or ‘aaah’; or
 - * It contains at least two words.
- To find out if a line is a *combined chords and lyrics* line, DECIBEL first extracts all characters between square brackets ([and]). If these characters form a chord, it removes all characters between each pair of square brackets. If the remaining line is a lyrics line, we can conclude that our original line was a combined chords and lyrics file.
- If a line does not satisfy any of these requirements, the parser assigns the line type *undefined*.

Segmentation

At this point, we have line types for each line in the tab. As a next step, our parser divides the tab file into segments. A *segment* is simply defined as the lines between lines of the empty line type. In a typical tab, verses and choruses are separated by empty lines and therefore are designated to their own segment.

System and chord extraction

Now we have our tab file divided into segments, and we know the line type of each of the lines in a segment. We want to extract the chord labels, but we only extract them from a suitable system. A *system* consists of all of lines that should be played or sung together at the same time. For example, in a typical chord sheet, we will find a lot of systems consisting of a *chord line* above a *lyrics line*, as the chords in the chord line should be played at the same time as the word in the lyrics line should be sung. The six different system types are listed in the first column of Table 6.1.

DECIBEL assigns each line to the largest possible system. Consider for example a segment consisting of three chord definition lines, followed by a chord line, six tablature lines and one lyrics line. Our parser finds only one system, consisting of a chord line, six tablature

System type	Chord Extraction
Combined chords and lyrics line	Parse chords from combined chords and lyrics line
Chord line	Parse chords from chord line
Chord line Exactly six tablature lines	Parse chords from chord line
Chord line Exactly six tablature lines 1 to 3 lyrics lines	Parse chords from chord line
Exactly six tablature lines	Parse chords from six tablature lines
Exactly six tablature lines 1 to 3 lyrics lines	Parse chords from six tablature lines

Table 6.1: System types and their chord extraction methods

lines and a lyrics line. In this case, the six tablature lines do not form a system on their own, as the chord line and lyrics line are also played and sung at the same time. This has to do with the very reason why we consider systems: we only extract one chord sequence from a system, to prevent duplicate chord sequences. This is particularly important in guitar tablatures, as they often have a chord line directly followed by six tab lines with exactly the same chord information. In that case, the parser ignores the six tab lines and extract the chords from the chord line only. For each of the systems, the chord extraction method is displayed in the second column of Table 6.1. We see that there are basically three chord extraction methods:

- **Parse chords from combined chords and lyrics line:** our parser extracts all characters between square brackets ([and]), so we have a list of strings, in which each string represents a chord. Then it parses each string by splitting it into a root note, chord type and optional bass note using regular expressions. For example, the chord string “D7/F#” has root note “D”, chord type “7” (dominant seventh chord) and bass note “F#”. Based on these three substrings, the parser now converts the chord string to Harte’s chord notation (Harte et al., 2005). It also saves the line number and character index of the start of the chord string.
- **Parse chords from chord line:** the parser extracts chord strings from the chord line by splitting on spaces and “[” characters. Then it parses each string and converts it to Harte’s chord notation, line numbers and character indices, as described above.
- **Parse chords from six tablature lines:** for each character index (x -value), the parser extracts the six characters that are notated above each other. If all characters are either a digit or an ‘x’, a chord is played at this position. The parser then derives the notes played at the same time based on these characters and their string - for example: if a digit ‘1’ appears on the B string line, the corresponding chord is a C; the digit ‘2’ on the B string corresponds to a C#. Finally it finds the nearest chord from our chord vocabulary (based on cosine distance). Again, it also saves line number (of the first tab line) and character index.




Figure 6.2: Recall of chords in our sample of 25 tab files.

6.1.2 Tab parsing evaluation

In this section, I describe the evaluation of the tab parsing algorithm. For the evaluation of the tab parsing algorithm, I used a random sample of 25 tabs from the data set, as listed in Appendix C. For these 25 tabs, I annotated the untimed chord sequence by hand. Then, I ran DECIBEL's tab parsing algorithm on these 25 tabs and compared the result to the annotation.

The results are shown in Figure 6.2, in which the height of each bar corresponds to the number of chords in the annotation of a tab file. The True Positives (TP) are the chords in the annotation which are also in the parsed results. In the figure, we see them as the blue part of the bar. The orange part of the bar corresponds to the False Negatives (FN): the chords that are in the annotation of a tab file, but not found by our parser. As we see in this figure, the parser finds (almost) all chords in most of the songs, as the orange part is very small. There were no chords in either the annotation or the parser output of the third, tenth and 23rd song. For the other songs, chords were missing in the tab output for the following reasons:

- The parser does not recognize chords that are glued after each other (without a separator) in the tab file;
- The parser does not recognize chords in a line with more text or strange symbols, as such a line is not classified as a chord line;
- The parser does not recognize chords in guitar tablature systems when there are no digits or 'x' in all six lines.

In some songs, text that was actually not a chord was misinterpreted as chords. This happened in one tab file which added some rhythm information by for example putting a 'Q' above a chord symbol for quarter note durations and an 'E' for eighth note durations. Our parser falsely detected a line of eighth notes as a sequence of E chords. In another line, chord vocabulary information was falsely detected as a chord line.

Still, in most songs we see that the tab parser performs quite well, as there are no missing chords in 15 of the 25 songs.

6.2 Jump Alignment

Having completed the tab parsing step, we have extracted the chord labels and their corresponding line and word numbers from the tab file. However, tab files retain no timing information, so we need an additional step to align the chord labels to the audio file. As we have seen in Section 1.3, there already exist four different algorithms by McVicar et al. (2011b) that incorporate tab information into a HMM-based system for audio chord estimation. To the best of our knowledge, these are the only algorithms that use tabs in audio ACE. The most promising of these four algorithms is Jump Alignment. In this section, I describe and evaluate our implementation of Jump Alignment.

Jump Alignment is based on a Hidden Markov Model (HMM). A HMM models the joint probability distribution $P(X, y|\Theta)$ over the feature vectors X and the chord labels y , where Θ are the parameters of the model. In Section 6.2.1, I describe the preprocessing steps that are used for extraction of the features X and the chord labels y . Section 6.2.2 gives an introduction to the HMM. The details on the Jump Alignment algorithm are treated in Section 6.2.3.

6.2.1 Preprocessing

First, the audio file needs to be **preprocessed**. For this purpose, we use the python package *librosa* (McFee et al., 2018). First, we convert the audio file to mono. Then, we use the HPSS function to separate the harmonic and percussive elements of the audio. Then, we extract chroma from the harmonic part, using constant-Q transform with a sampling rate of 22050 and a hop length of 256 samples. Now we have chroma features for each sample, but we expect that the great majority of chord changes occurs on a beat. Therefore, we beat-synchronize the features: we run a beat-extraction function on the percussive part of the audio and average the chroma features between the consecutive beat positions. The chord annotations need to be beat-synchronized as well. We do this by taking the most prevalent chord label between beats. Each mean feature vector with the corresponding beat-synchronized chord label is regarded as one frame. Now we have the feature vectors X and chord labels y for each song, which we feed to our HMM.

6.2.2 Hidden Markov Model

As we have seen in Section 4.1.2, a HMM models the joint probability distribution $P(X, y|\Theta)$ over the feature vectors X and the chord labels y , where Θ are the parameters of the model. The HMM is an extension of the Markov chain in which the states are not

directly observable (but hidden). In this section, I give a summary of the Hidden Markov Model, based on the tutorial by Rabiner (1989) and Section 5.3.2 from Müller (2015).

In a **Markov chain**, we make the assumption that the probability from the current state s_t to the next state s_{t+1} only depends on the current state, and not on any of the previous states. Formally:

$$P(Y_{t+1} = s_{t+1} | Y_1 = s_1, Y_2 = s_2, \dots, Y_t = s_t) = P(Y_{t+1} = s_{t+1} | Y_t = s_t) \quad (6.1)$$

This assumption is called the **Markov property**. In our implementation, states are represented by chord labels. Any chord label is an element from the discrete set of chords from our chord vocabulary. Although we could use any chord vocabulary, we use the vocabulary with all major and minor chords and the no-chord symbol. Our vocabulary thus consists of 25 items. Thanks to the Markov property, we can efficiently model the probability that chord j occurs after chord i by the **transition probability** $P_{tr}(y_t = j | y_{t-1} = i)$. Given a chord vocabulary of 25 items, we can model all transition probabilities in $25 \cdot 25 = 625$ probability values, that can be efficiently stored in the matrix \mathbf{P}_{tr} .

The first chord of a chord progression cannot be computed from its previous chord. This information is specified by the **initial state probabilities**. There are 25 initial state probability values, stored in the vector \mathbf{P}_{ini} , in which $P_{ini}(y_1 = c)$ is the probability that the first chord is chord c .

Based on the Markov chain, we could now compute the probability for a given chord label given a sequence of chord labels, or compute the probability of any chord label sequence. However, in a ACE scenario, we cannot directly observe the chord labels, but we do observe a sequence of chroma features that is related to the chord labels. We thus need to expand our model, by adding **emission probabilities**. The resulting model is a Hidden Markov Model. The **emission probability** $P_{obs}(X_t | y_t)$ equals the probability of an observation (chroma feature vector) X given a chord label y at time t . In contrast to transition and initial state probabilities, the emission probabilities cannot be represented by a matrix, as chroma features are continuous. Instead, they are modeled by continuous probability density functions: the observation probability distribution for each of the 25 chord labels a_i in the vocabulary as a 12-dimensional Gaussian with 12-dimensional **mean vector** μ_i and 12×12 **covariance matrix** Σ_i . We can now find our emission probabilities for any observed chroma feature vector x_t and any chord label a_i , using the Gaussian with parameters μ_i and Σ_i :

$$P_{obs}(x_t | \mu_i, \Sigma_i) = \frac{1}{(2\pi)^6 |\Sigma_i|^{1/2}} \exp\left[-\frac{1}{2}(x_t - \mu_i)^T \Sigma_i^{-1} (x_t - \mu_i)\right] \quad (6.2)$$

We now have seen that a HMM is parametrized by transition, initial state and emission probabilities. We define Θ as the set of these parameters:

$$\Theta = \{\mathbf{P}_{tr}, \mathbf{P}_{ini}, \{\mu_i\}_{i=1}^{25}, \{\Sigma_i\}_{i=1}^{25}\} \quad (6.3)$$

There exist different ways of estimating Θ . We adopt the machine learning approach, that is: we **train** the probabilities on a subset of our data set. In our case, the training set

is a 100-song random sample of our data set. The initial state and transition probability matrices are calculated simply by frequency counting and normalization of all chord labels in the training set. For calculating μ_i and Σ_i , we take the set of all feature vectors that are labeled with chord i in our data set, and then calculate the mean vector and covariance matrix respectively.

Now we know all elements of the HMM, we can represent the joint representation for the chroma feature vectors X and chord labels y mathematically in the following equation:

$$P(X, y|\Theta) = P_{ini}(y_1|\Theta) \cdot P_{obs}(x_1|y_1, \Theta) \cdot \prod_{t=2}^{|y|} P_{tr}(y_t|t_{t-1}, \Theta) \cdot P_{obs}(x_t|y_t, \Theta) \quad (6.4)$$

HMMs can be used to solve three basic problems, as pointed out in Rabiner (1989). One of these problems, which is called the **uncovering problem** or **decoding problem**, is the relevant problem for our purpose: automatic chord estimation. The goal of the uncovering problem is to find the state (chord) sequence that best explains the observed (chroma) feature sequence. Formally: we want to find the chord label sequence y^* where $y^* = \operatorname{argmax}_y P(y|X, \Theta)$. Note that this is equivalent to finding $y^* = \operatorname{argmax}_y P(y, X|\Theta)$.

Note that the optimal chord sequence for an observation sequence of n elements starts with the optimal chord sequence for the observation sequence of the first $n - 1$ elements. The **Viterbi algorithm** efficiently solves the uncovering problem by exploiting this property. The pseudocode of Viterbi is given in Algorithm 2. In each iteration, the algorithm computes the indices i_1, i_2, \dots, i_t to the optimal chord sequence for X_1, X_2, \dots, X_t (i.e. the first t steps of the observation sequence X) ending with chord label y_{i_t} . The result is then stored in the matrices V and B : in V , the algorithm keeps the likelihood of the current path, while the chord index is stored in B . When calculating a chord sequence of length n , the likelihoods of the optimal chord sequences of length $n - 1$ can be efficiently retrieved by reading the corresponding value from matrix V . In the termination phase, we can read the final chord index i_T by maximizing over the last column of B and find the full optimal chord sequence by following the backpointers in B back in time. We can also read the likelihood of the full optimal chord sequence from $V[i_T, T]$.

6.2.3 Jump Alignment

Jump Alignment is an extension to the HMM described in the previous section, which utilizes the chords that are parsed from tabs. Following McVicar et al. (2011b), we refer to these chords parsed from tab files as Untimed Chord Sequences (UCSs). Compared to the original HMM, in the Jump Alignment algorithm the state space and transition probabilities are altered in such a way that it can align the UCSs to audio, while allowing for jumps to the start of other lines. I will explain how this works in the remainder of this section.

Let us assume that we have an audio file and tab file from the same song. Let $\mathbf{e} = e_1, e_2, \dots, e_L$ be the UCS we extracted from the tab file, applying the parsing algorithm

Algorithm 2 Pseudocode of the Viterbi algorithm

```

1: function VITERBI(Observation Sequence  $X_1, X_2, \dots, X_T$ , HMM specified by  $\Theta$ ,
   Chord alphabet  $y_1, y_2, \dots, y_{25}$ )
2:   Initialize Viterbi Matrix  $V$  of size  $25 \times T$  ▷ Initialization
3:   Initialize Backpointer Matrix  $B$  of size  $25 \times (T - 1)$ .
4:   for all  $i \in [1, 25]$  do
5:      $V[i, 1] = \mathbf{P}_{\text{ini}}[y_i] \cdot P_{\text{obs}}(y_i | X_1)$ 
6:   end for
7:   for all  $t \in [2, T]$  do ▷ Recursion step
8:     for all  $i \in [1, 25]$  do
9:        $V[i, t] = \max_{j=1}^{25} V[j, t-1] \cdot \mathbf{P}_{\text{tr}}[y_j, y_i] \cdot P_{\text{obs}}(y_j | X_t)$ 
10:       $B[i, t-1] = \operatorname{argmax}_{j=1}^{25} V[j, t-1] \cdot \mathbf{P}_{\text{tr}}[y_j, y_i]$ 
11:     end for
12:   end for
13:    $i_T = \operatorname{argmax}_{j=1}^{25} V[j, T]$  ▷ Termination
14:   for all  $t \in [T-1, 1]$  do
15:      $i_t = B[i_{t+1}, t]$ 
16:   end for
17:   return Chord Sequence  $y_{i_1}, y_{i_2}, \dots, y_{i_T}$  and Likelihood Value  $V[i_T, T]$ 
18: end function

```

described in Section 6.1. Because we want an alignment between the audio file and this UCS \mathbf{e} , we define a new set of hidden states: $Y' = \{1, 2, \dots, L\}$ corresponds to the indices of the chords in \mathbf{e} . For example, given a (very small) chord vocabulary $[C, D, E]$ and the UCS $[C, E, D, E, D, C, E]$, our hidden state set Y' would be $[1, 3, 2, 3, 2, 1, 3]$. Note that different hidden states can now refer to different occurrences of the same chord, in contrast to our original state space.

We also need to alter the transition probability distribution P'_{tr} , as we only want to allow chord sequences that are alignments of \mathbf{e} , our Untimed Chord Sequence. However, McVicar et al. (2011b) showed that exact alignment of \mathbf{e} to audio results in a decrease in performance compared to the original HMM, because repetition cues (like “play this verse twice”) are not interpreted by the tab parser. In addition, some tabs do not specify the chords of all verses. Therefore, Jump Alignment allows jumps from the end of a line in the tab file to the beginning of any line (the current one, of any of the previous or subsequent lines). The new transition probability distribution is expressed as:

$$P'_{tr}(j|i, \Theta, e) = \begin{cases} \frac{1}{Z_i} P_{tr}(e_i, e_j) & \text{if } j \in \{i, i+1\}, \\ \frac{p_f}{Z_i} P_{tr}(e_i, e_j) & \text{if } i < j \text{ and } i \text{ is the end and } j \text{ is the beginning of a line,} \\ \frac{p_b}{Z_i} P_{tr}(e_i, e_j) & \text{if } i > j \text{ and } i \text{ is the end and } j \text{ is the beginning of a line,} \\ 0 & \text{otherwise.} \end{cases} \quad (6.5)$$

Z_i is a normalization factor, which is used to re-normalize P'_{tr} so that the transition probabilities $P'_{tr}(y_j|y_i, \Theta, e)$ sum to one. p_f is the probability of jumping forward when this is allowed. Likewise, p_b is the probability of a jump backwards, provided that this is allowed.

	WCSR
Average CSR of all tabs	0.738
Best log-likelihood of all tabs	0.770
Best CSR of all tabs	0.778

Table 6.2: WCSR of all songs, with different tab file selection methods

The forward probability p_f and backward probability p_b need to be specified. We estimate the optimal p_f and p_b as follows: first, we manually select a subset of 22 songs for which we are sure that they are in the right key (so we do not need to consider transpositions). Appendix D lists these songs. Then, we run the Jump Alignment algorithm with all combinations of forward and backward probabilities between 0 and 1 in steps of 0.05 and evaluate the resulting chord label sequences in terms of CSR. We will use the p_f and p_b with the highest average CSR.

Figure 6.3 shows the average CSR for our 22-song subset, given a (p_f, p_b) parameter setting. We see that positive forward or backward probabilities improve CSR compared to a strict alignment (in which $p_f = 0$ and $p_b = 0$): particularly non-zero backward probabilities yield considerably better chord label sequences than p_b 's of zero. However, if the forward and backward probabilities become too large, the average CSR decreases as the importance of the order of lines in the tab file decreases. We find an optimal parameter setting of $p_f = 0.05$ and $p_b = 0.05$. This is the parameter setting I use in my implementation of Jump Alignment.

Tabs are often notated in a transposed key compared to the original audio file, because some keys are easier to play on the guitar than others. In order to correct for such “simplified” tabs, Jump Alignment considers the tab files in all 12 transpositions and chooses the transposition with the highest likelihood.

6.3 Tab file selection

The performance of the jump alignment algorithm is dependent on the quality of the tab file. Following McVicar et al. (2011b), we select for each song the tab file for which the log-likelihood was the highest. The results for our full 200 song dataset are shown in Table 6.2. If we would take the average CSR of all tabs for a song, the WCSR is 73.8%. The upper limit for the selection method is a WCSR of 77.8%: this is what we would get if we selected the best tab. However, we do not know yet which tab file will be the best but estimate this based on the log-likelihood. When selecting the tab file with the highest log-likelihood, we reach a WCSR of 77.0%, which is very close to the upper limit.




Figure 6.3: Average CSR for 22-song subset given p_f and p_b . We find an optimal parameter setting of $p_f = 0.05$ and $p_b = 0.05$.

6.4 Conclusion

In this section, we have examined the last of the three representation-specific subsystems of DECIBEL: the tab subsystem. This subsystem consists of a parser that extracts Untimed Chord Sequences (UCSs) from guitar tablature or chord sheets, thereby retaining line information. This is then input for the Jump Alignment algorithm, which aligns the UCSs to the audio recording. When selecting the estimated best tabs for each song, based on log-likelihood, this subsystem reaches a WCSR of 77.0% on the Isophonics data set.

Chapter 7

Data fusion

DECIBEL estimates chord label sequences from different music representations, i.e. audio, MIDI and tab files, as we have seen in the previous three chapters. This results in a set of chord label sequences for each song in our data set. This set of chord labels forms a rich harmonic representation that is already interesting in itself, as we will see in Section 7.1. However, in order to answer our research question we need to combine these chord label sequences into one final sequence (and compare the resulting chord sequence to the sequence obtained by using only an audio ACE method). DECIBEL achieves this using a data fusion step. In this chapter, we motivate the chosen data fusion method by a literature study and an evaluation of various methods.

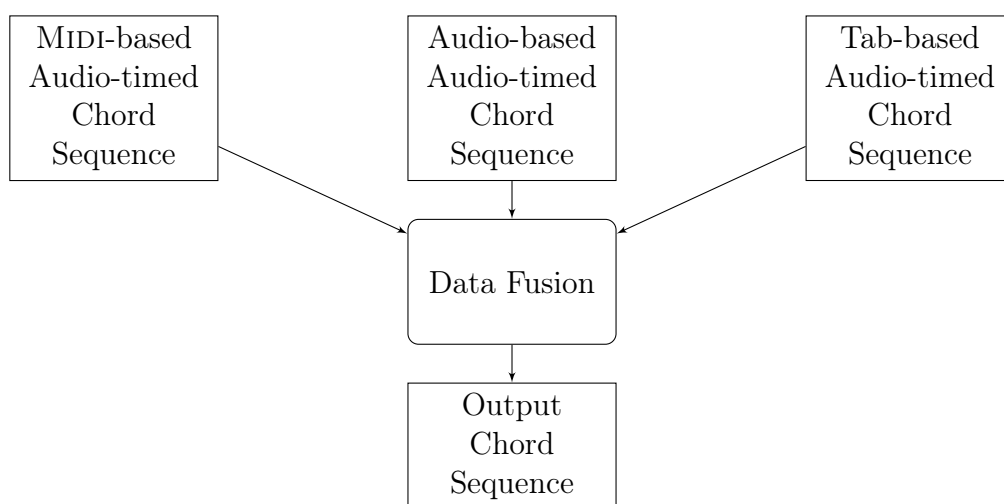


Figure 7.1: Diagram of DECIBEL’s data fusion subsystem. The set of MIDI- audio- and tab-based audio-timed chord label sequences forms a rich harmonic representation. Our data fusion combines all chord label sequences for a song into a single output chord label sequence.

7.1 A rich harmonic representation

In the previous chapters we have seen how DECIBEL’s subsystems estimate chord label sequences from audio, MIDI and tab representations of a song. These three representations store musical content of the same song in fundamentally different ways: as we have seen in Section 2.5, audio files are obtained by digitizing the waveform; MIDI files are a series of note on and note off events; whereas tab files show the guitar fingering or chord labels, aligned to the lyrics. The timing information in audio is represented in seconds, based on the underlying performance; in MIDI files, timing is measured in ticks; in tabs, timing information is missing. Also, the source of information differs: audio files are recordings of a performance; MIDI files are either score-based or transcriptions of a recording; tabs can be considered as (untimed) chord label sequences, manually annotated by music enthusiasts.

However, the chord label estimation and synchronization steps performed by DECIBEL’s subsystems transform these three heterogeneous representations into a homogeneous format, i.e. as a series of $\langle \text{start time, end time, chord label} \rangle$ 3-tuples. The combination of the set of chord label sequences from each of the representations of a song results in a rich harmonic representation.

This rich harmonic representation is a very interesting by-product of the DECIBEL system, as it allows for cross-version analysis, i.e. comparing analysis results from different representations. The chord label sequences can easily be visualized, which makes it very easy to see the consistencies and inconsistencies in chord labels between different representations. As identified earlier in research by Ewert et al. (2012) and Konz et al. (2013), having such a unified view of different analysis results deepens the understanding of both the algorithm’s behavior and the properties of the underlying musical material.

Consider for example the visualization of the chord label sequence for the Beatles song *Golden Slumbers* in Figure 7.2. In our data set, we have matched three MIDI files and eight tab files to the audio version of the song. CASSETTE, our MIDI ACE method, analyzes the chords both on a bar and a beat level, so we have six chord label sequences based on the three MIDI files. In combination with the eight chord label sequences based on the tab files and a single analysis by the Chordify algorithm based on the audio file, we have $6 + 8 + 1 = 15$ chord label sequences. Each of these estimated chord sequences, as well as the ground truth, is visually represented by a horizontal bar in which the color represent the chord label and time in seconds can be read from the position on the x -axis. In this example, we observe there are some parts of the song for which (most of) the 15 estimated chord sequences agree on the labels, e.g. the D minor chord starting after 11 seconds; the G major chord sounding immediately thereafter or the C major chord that starts at 32 seconds. However, in other musical passages there is some disagreement. If there is a single deviating chord sequence estimation, this is typically either due to some error in the MIDI or tab file or due to an error in the estimation method. For example, when estimating the chord sounding in the 38th second from our third MIDI file, using CASSETTE on the bar level, we falsely find a F major chord instead of a C major chord. In some passages, we observe multiple “clusters” of possible chord labels, i.e. different perspectives on the harmonic content of a musical passage, each of them supported by

multiple chord label estimations. For example, we can consider the start of the song either as a long A minor chord or as a sequence of A minor and C major chords. Also, there exist different perspectives on the chord in the 23rd second (is it a C major, E major, F minor or something else?) and the modality of the A (major or minor) chord starting after 62 seconds.

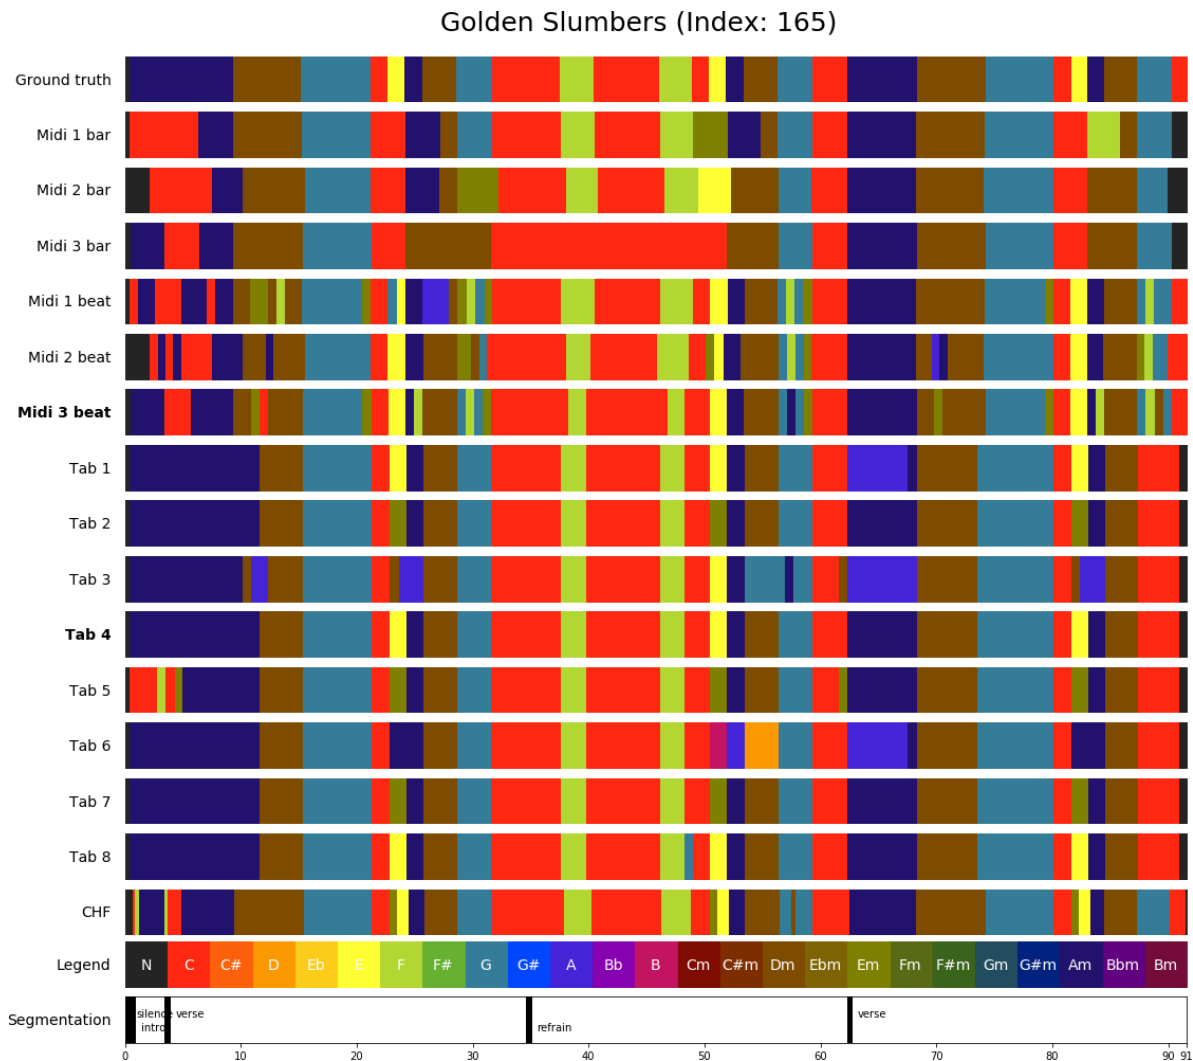


Figure 7.2: Multiple perspectives on the Beatles song *Golden Slumbers*: we can consider the start of the song as a long A minor chord or as a sequence of A minor and C major chords and there are multiple opinions on e.g. the modality of the chord starting after 62 seconds.

MIDI and tab files implicitly incorporate musical knowledge, as they are (in the popular music genre) typically transcriptions made by music enthusiasts. Therefore, the comparison of analysis results from different music representations reveals passages in the music for which there are multiple possible perspectives. Studying these passages can give us a deeper understanding of subjectivity issues in the musical material.

In this section, we have seen that the chord label sequences estimated on different music representations of the same song together form a rich harmonic representation. Visualization of this harmonic representation can give valuable insights in different perspectives on chord labels within a specific segment of the song. This makes large-scale cross-version

analysis feasible. Moreover, it bridges the gap between technical, audio-oriented, musical signal processing and non-technical, score-oriented musicology. A detailed analysis on the different harmonic perspectives on the Beatles and Queen songs from our data set is beyond the scope of this research, but would be an interesting direction of future work which would provide better insights in the subjectivity issues in harmonic analysis of popular music.

7.2 Related Work on data fusion

In the remainder of this chapter, we study data fusion methods to combine chord label sequence obtained from DECIBEL’s audio, MIDI and tab subsystems into a single output sequence. The integration of heterogeneous output of multiple ACE algorithms using data fusion is a new idea, which was recently proposed by Koops et al. (2016). In this study, the authors experiment with three different techniques to combine chord sequence estimates from different sources (the MIREX 2013 ACE submissions, applied to the Billboard data set) into one final output sequence for each song. They show that their proposed data fusion method yields the best results in terms of WCSR. Also, they show that the output sequence found by their data fusion method is an improvement to the best scoring team, with an increase between 3.6 percentage point and 5.4 percentage point compared to the best team.

The three techniques used for combining chord label sequences are random picking (RND), majority voting (MV) and data fusion (DF). In all three techniques, the system proposed in Koops et al. (2016) first samples the chord label sequence from each source in segments of 10 milliseconds. In the RND technique, the output chord label for each segment is found by picking the corresponding label from a randomly chosen source. For MV, the output chord label for each segment is the most frequent label of all sources. If multiple labels are most frequent, the output label is picked randomly from these most frequent labels. The DF technique takes three things into account when integrating the output of multiple ACE algorithms:

- The accuracy of sources;
- The probabilities of the values provided by the sources; and
- The probability of dependency between sources.

The **source accuracy** $A(S_i)$ is the arithmetic mean of the probabilities of all chord labels a source S_i provides. Initially, these chord label probabilities may be based on a frequency count. A source with a high source accuracy inherently has a lot of chord labels that agree with other sources, and therefore can be considered more trustworthy. Assuming that the sources are independent, the source accuracy is the probability that a source provides the appropriate chord. On the other hand, the probability that a source provides an inappropriate chord can be computed with $\frac{1-A(S_i)}{n}$, given n possible inappropriate chord labels. The **source vote count** $VS(S_i)$ combines these probabilities and is computed as

follows: $VS(S_i) = \ln \frac{n \cdot A(S_i)}{1 - A(S_i)}$. Chord labels of sources with higher source vote counts are more likely to be picked than labels of sources with low source vote counts, as we will see shortly.

The goal of the data fusion method is to determine the most likely chord label for each segment, given a number of sources. To this end, the data fusion method computes a **chord label vote count** $VC(\mathcal{L})$ for each chord label \mathcal{L} . The chord label vote count is dependent on both the number of sources with a matching chord label and the source vote count of these sources, as expressed in the formula $VC(\mathcal{L}) = \sum_{\sigma \in S^{\mathcal{L}}} VS(\sigma)$, in which $S^{\mathcal{L}}$ is the set of sources with chord label \mathcal{L} .

From the label vote counts, we can compute **chord label probabilities** for each chord label \mathcal{L} by $P(\mathcal{L}) = \frac{\exp(VC(\mathcal{L}))}{\sum_{l \in D} \exp(VC(l))}$. As we divide the chord label vote count of our chord label by the chord label vote count of all possible chord labels D , this results in a probability value between 0 and 1.

The data fusion method proposed by Koops et al. (2016) also takes the dependency between sources into account, following the intuition that sources that are dependent of each other share many uncommon chord labels. Based on the number of shared uncommon chords with other sources, the authors compute a **source dependency** weight $\mathcal{I}(S_i, \mathcal{L})$, which represents the probability that a source S_i provides a chord label \mathcal{L} independently.

We now have seen the definitions of chord label probabilities, source accuracy and source dependency. Note that they are all defined in terms of each other. As a solution to this paradoxical situation, Koops et al. (2016) propose to initialize the chord label probabilities with equal probabilities and iteratively compute source dependency, chord label probabilities and source accuracy until convergence is reached. Finally, their data fusion method selects for each segment the chord label with the highest chord label probability.

In their experiments, the authors compare the performance of DF to RND and MV for two subsets of the Billboard data set. Each of these three chord sequence combination methods is computed on all twelve MIREX 2013 ACE submissions for both Billboard subsets. The Billboard annotations serve as ground truth. Their results show that the DF method significantly outperforms both the RND and MV methods as well as each of the individual submissions. We can conclude that combining chord label sequences from multiple sources using data fusion is a promising strategy to improve ACE.

7.3 Data fusion experiments

In the previous section, we have seen that the heterogeneous output of multiple ACE systems can be integrated into a single output chord sequence, using data fusion. Although the results obtained in previous work are solely based on audio data, the data fusion method can also be applied to chord label sequences that are based on symbolic data. This is exactly what DECIBEL’s data fusion subsystem does: the subsystem uses the chord sequences from audio, MIDI and tabs and integrates them in a final output sequence. In

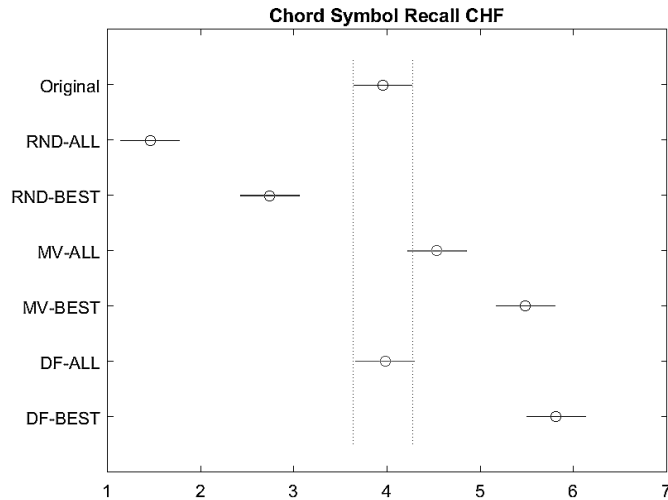


Figure 7.3: Visual representation of the differences in terms of Chord Symbol Recall between different data fusion methods, for the CHF audio ACE algorithm. For each of the horizontal lines that do **not** overlap in the y -direction, the difference in CSR between the corresponding data fusion methods is significant. For example, DF-BEST is significantly better than DF-ALL, but the difference between DF-BEST and MV-BEST is not significant.

contrast to earlier work, in our case the number of sources differs per song. That is why I experiment with a selection strategy that elects only one source per representation. In this section, I will explain DECIBEL’s data fusion implementations.

In this section, I compare two selection strategies in combination with three different integration methods. The two selection strategies are ALL and BEST: ALL takes the chord sequences of all tabs and MIDI files as sources. BEST only uses the sources of the expected best tab and MIDI file for each song. This way, the integration method is always applied to only three sources.

The three integration methods are based on the earlier work by Koops et al. (2016). Following their work, we first sample each input chord sequence in 10 millisecond steps. Then we integrate the sources (selected using one of the two selection strategies) with either random picking (RND), majority voting (MV) or data fusion (DF). The implementations of RND and MV are unchanged compared to earlier work. For the DF integration method, I made two alterations: (1) I omitted the source dependency weight, because our MIDI and tab subsystems only use the major-minor chord vocabulary. Therefore, there is no matter of rare chords as in earlier work by Koops et al. (2016). (2) I empirically tested the number of iterations needed for reaching a fixed point, and henceforth always terminate the DF algorithm after this number of iterations, without checking if a fixed point was reached.

7.4 Data fusion results

We now compare the performance of each of the six combinations of integration methods and selection strategies (i.e. RND-ALL, RND-BEST, MV-ALL, MV-BEST, DF-ALL and DF-BEST), applied to each of our seven audio ACE systems. Friedman tests (Friedman, 1937) for each of these seven audio ACE show that the integration methods and selection strategies give significantly different results in terms of CSR. Consequently, I perform a Tukey’s Honest Significant Difference post-hoc test (Tukey, 1949) to identify which integration-selection combinations are significantly different. The results can be found in the Appendix E and are visually represented in Figures 7.3 and 7.4. In these figures, we can easily observe which differences are significant because the corresponding horizontal lines do not overlap in the y -direction. Consider for example Figure 7.3. Here we see that both RND methods are significantly worse than the original audio ACE system (CHF). There is no significant difference between CHF and DF-ALL. MV-ALL is better than CHF, but the difference is not found significant by Tukey’s HSD test. MV-BEST and DF-BEST perform significantly better than the original CHF system, and DF-BEST turns out to be the winner.

Let us now consider the test results for all seven state-of-the-art audio ACE systems in combination with all two selection strategies and all three integration methods. We can make the following three observations:

As a first observation, the random picking integration method always performs worse than the original audio ACE system, regardless of the chosen selection strategy. This is consistent with the findings in earlier work by Koops et al. (2016). From this observation, we can conclude that it is interesting to take the agreements and differences between sources into account.

Secondly, the BEST selection method always performs better than the ALL selection method. The explanation for this is following: given a poor MIDI or tab file, DECIBEL’s MIDI or tab subsystem will find a poor chord label sequence. In the ALL selection method, this chord label sequence will, to a greater or lesser extent, be integrated in the final output sequence, which deteriorates the output sequence’s quality, whereas the BEST will (hopefully) ignore this poor chord label sequence. We observe that the difference between BEST and ALL is even larger for the DF integration method than for MV. A possible explanation for this observation could be that chord label sequences from tabs are often undersegmented, for example because chord changes in instrumental parts or very short chords are often not detected. This is a typical property for tabs and therefore can occur at multiple (independent) tabs of the same song. As a result, suboptimal chord sequences from tabs may get a too high source accuracy in DF-ALL, and consequently have a large influence on the final output chord sequence. In DF-BEST, we only compare the best tab to the best audio and the best MIDI file. If the best tab is undersegmented, it will probably get a low source accuracy as chord sequences from MIDI and audio have less undersegmentation issues. The difference between MV-ALL and MV-BEST is smaller, because MV does not use source accuracies.

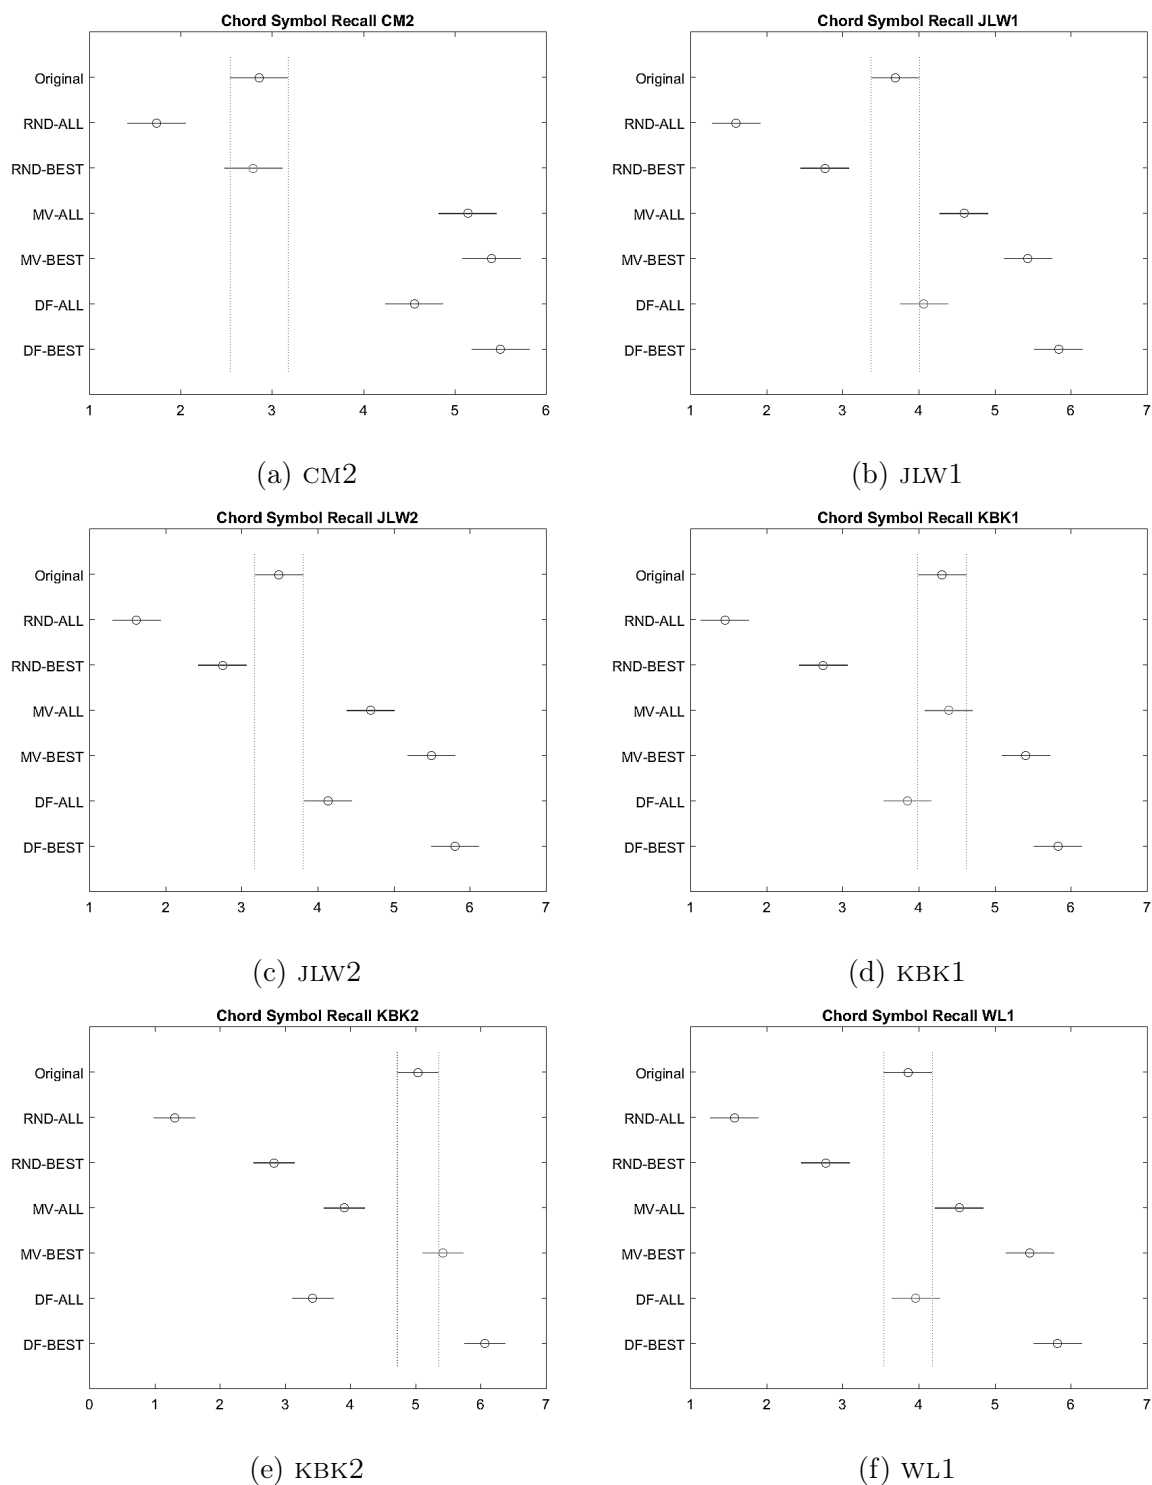


Figure 7.4: Visual representation of the differences in terms of Chord Symbol Recall between different data fusion methods, for each of the MIREX 2017 ACE submissions. For each of the horizontal lines that do **not** overlap in the y -direction, the difference in CSR between the corresponding data fusion methods is significant.

Our third and final observation is that DF-BEST performs better than MV-BEST, although the difference is not always significant. However, DF-BEST is the only method that is always significantly better than the original audio ACE system. Even considering KBK2, the best-performing audio ACE system of our seven state-of-the-art systems, we see that DF-BEST improves it significantly.

When comparing the oversegmentation measure for our integration methods, we can do another observation: RAND and MV tend to oversegment, in contrast to DF. This becomes clear when looking at the chord sequences for the song *Things We Said Today* in Figure 7.5. Note that both RND chord sequences have way too many chord changes, resulting in a low oversegmentation score. Both MV chord sequences are oversegmented as well, albeit to a lower degree. On the other hand, we see no oversegmentation issues in the DF chord label sequence.




Figure 7.5: Chord label sequences for the Beatles song *Things We Said Today*. CSR = Chord Symbol Recall; OvS = oversegmentation; UnS = undersegmentation; and Seg = segmentation. Note that RAND and MV tend to oversegment, in contrast to DF.

From these observations, we can conclude that DF-BEST is the best selection-integration combination. Table 7.1 shows a comparison between the WCSR of each of the seven audio ACE systems and DF-BEST applied to each system. The table shows that using DF-BEST improves ACE WCSR on average by 3.7%.

Audio ACE system	Audio WCSR	DF-BEST WCSR	Improvement
CHF	.820	.848	2.8%
CM2	.752	.817	6.5%
JLW1	.785	.827	4.2%
JLW2	.779	.826	4.7%
KBK1	.822	.852	3.0%
KBK2	.867	.875	0.8%
WL1	.793	.834	4.1%

Table 7.1: WCSR of audio ACE systems and DF-BEST. Using DF-BEST improves ACE WCSR on average by 3.7%.

Finally, let us look at the song-wise performance of DF-BEST compared to the original audio file. Figure 7.6 shows the difference between the CSR of the chord label sequence obtained by DF-BEST and the CSR of the chord label sequence found by the original audio ACE system. The song keys are given on the y -axis, and the length of the horizontal bars correspond to the difference in CSR. If a bar is located on the right side of $x = 0$, then the chord label sequence found by DF-BEST is better than the sequence found by the original audio ACE system. We see that this is the case for the vast majority of songs in our dataset. On the other hand, if a bar is located on the left side of $x = 0$, DF-BEST performs worse than the original audio ACE system. From all seven plots, we see one song for which DF-BEST performs considerably worse. This is song number 174, *Let it be* by The Beatles. This song has an issue similar to *Don't pass me by*, as we saw in Figure 4.5: the audio file I used was shifted compared to the audio file that was used in the MIREX competition. That is why the chord label sequence found by DF-BEST is shifted and therefore is not consistent with the Isophonics annotations any more. We see another ‘peak’ to the left at song 196, *We will rock you* by Queen, for the CHF algorithm. The reason for this difference is that CHF estimates the chord label sequence for this song quite well, as shown in Figure 7.7. However, the best found MIDI and tab file find very different chord labels for the first 90 seconds of the song and therefore the result found by DF-BEST is based on the tab file, which apparently had the highest source accuracy. In practice, this song starts with a lot of percussion and little harmonic content. Therefore, it is not so surprising that there are multiple views on the chord labels for this song.

In this section we have seen that DF-BEST is the best combination of the selection strategy and integration method. DF-BEST performs significantly better than the original audio algorithm in terms of WCSR and does not suffer from oversegmentation.

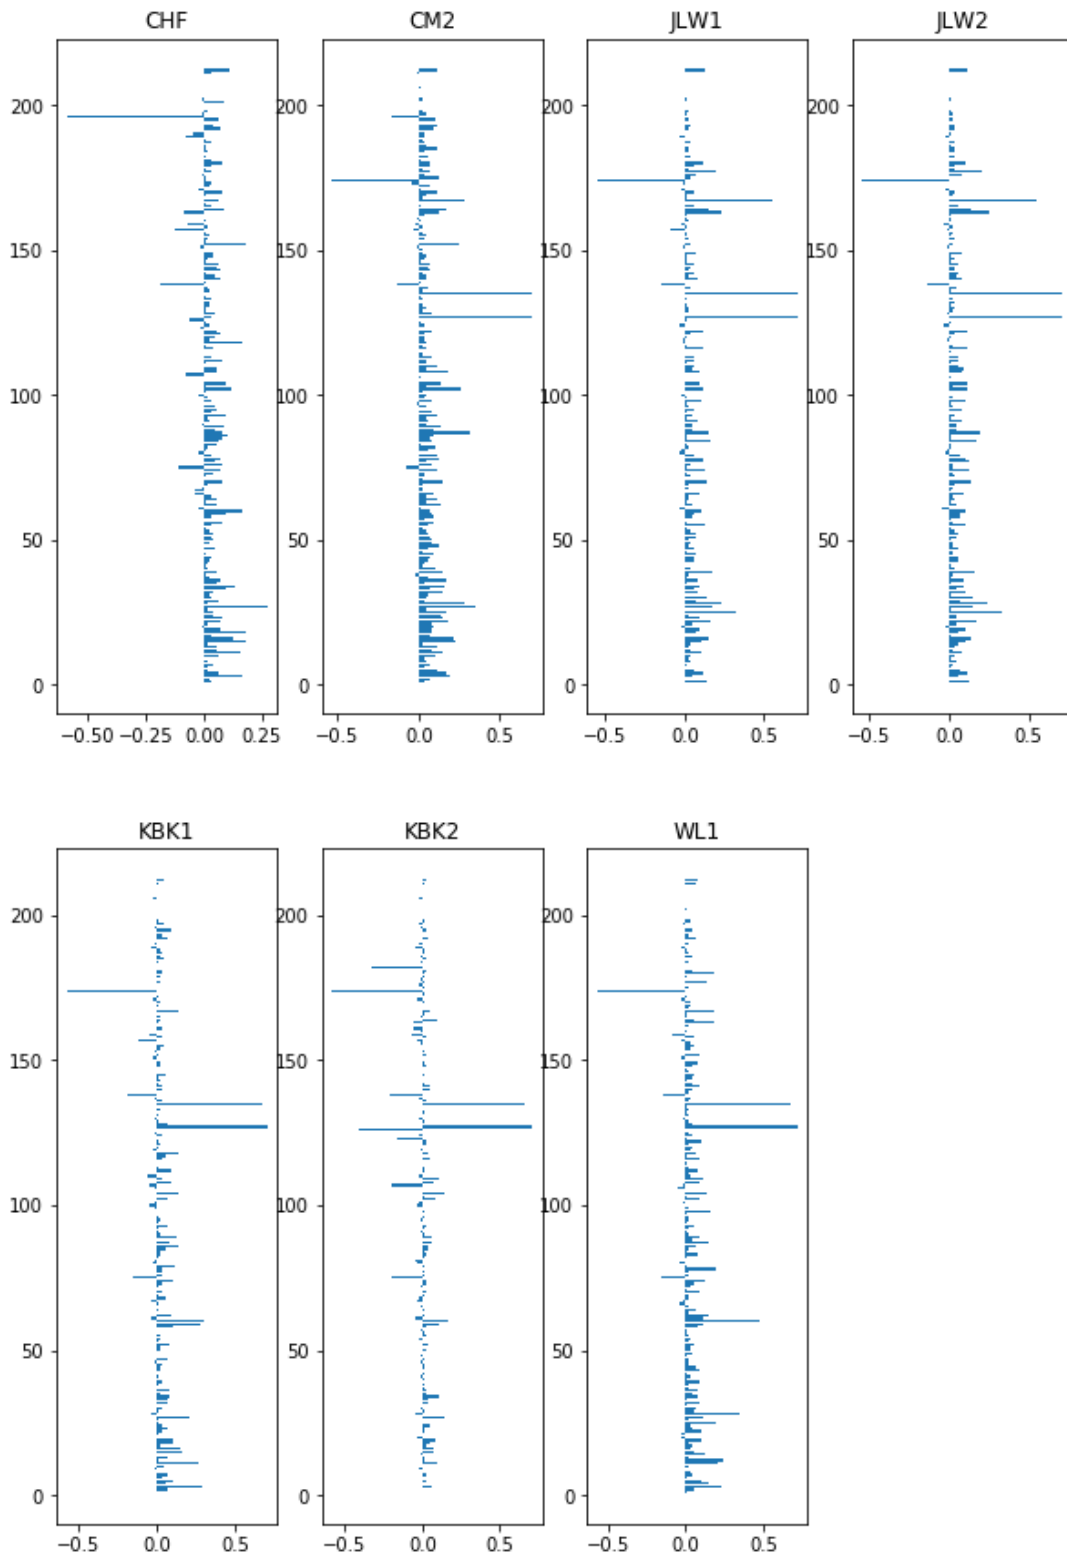


Figure 7.6: CSR of DF-BEST compared to the CSR of the original audio ACE system for each song in the dataset. The difference between the CSR of DF-BEST and the CSR of the original audio ACE system is given on the x -axis; the y -axis gives the index of the song in our data set. Horizontal lines on the right side of $x = 0$ correspond to songs for which DF-BEST improves the original system in terms of CSR, while the horizontal lines on the left side of $x = 0$ correspond to songs for which the original ACE system produced better results in terms of CSR.




Figure 7.7: Chord label sequences for the Queen song *We will rock you*. Note: CSR = Chord Symbol Recall; OvS = oversegmentation; UnS = undersegmentation; and Seg = segmentation.

Chapter 8

Conclusion

The goal of this thesis was to test whether exploiting MIDIs and tabs improves Automatic Chord Estimation (ACE). The ACE task is concerned with estimating chords in audio recordings or symbolic representations of music and has applications in both music performance and in Music Information Retrieval (MIR). Many systems for audio ACE exist, but ACE is not yet a solved problem: as we have identified in Section 1.1, existing methods suffer from stagnation and overfitting to subjective annotations. There is need for a new strategy that overcomes existing stagnation in ACE without further overfitting to existing (subjective) data sets. An improved system for ACE would have positive consequences for millions of performers and MIR systems for e.g. cover song identification and genre classification.

For this purpose, I created a new system, DECIBEL, that (1) aligns MIDI and tab files to audio recordings and uses representation-specific chord estimation techniques for both symbolic and audio formats to estimate chord sequences for each file; and (2) integrates the resulting heterogeneous chord sequences into one final output sequence. In this concluding chapter I will give an overview of my contributions and some directions for future work.

8.1 Contributions

The main contributions of this research are twofold: I designed a system that automatically extracts a rich harmonic representation and used this representation to improve state-of-the-art audio ACE systems. In addition, I extended the existing Isophonics Reference Annotations with a large set of MIDI and tab files, which I manually matched to the annotations. As a final contribution, I made the implementation of my system available on GitHub.

8.1.1 Rich harmonic representation

As we have seen in Chapters 4, 5 and 6, DECIBEL has three subsystems that estimate chord label sequences from audio, MIDI and tab files, using representation-specific ACE and alignment techniques. The resulting chord label sequences form a rich harmonic representation that is a contribution by itself, as it can easily be visualized and enables large-scale cross-version analysis of popular music, giving more insights on subjectivity issues. The potential use of this rich harmonic representation is described in Section 7.1.

8.1.2 Improvement of state-of-the-art audio ACE systems

In Chapter 7, I described my experiments on data fusion. DECIBEL uses data fusion methods to combine the chord label sequences from the aforementioned rich harmonic representation into a single output sequence. The method DF-BEST turned out to be the best-performing method. Using this data fusion method, DECIBEL significantly improves each of the seven original state-of-the-art audio ACE systems. The average WCSR improvement is as much as 3.7%. Supported by statistical tests, I can confirm the hypothesis that audio ACE can be improved by integrating symbolic music formats.

Although DECIBEL needs a data set of symbolic formats and has a longer calculation time than the original audio ACE systems, it also has two major advantages beyond the WCSR improvement. As a first advantage, by using MIDI and tab files that are scraped from the Internet, DECIBEL implicitly incorporates musical knowledge: MIDI and tab files can be considered as human-made transcriptions of musical notes and chords. As a second advantage, DECIBEL requires a minimum amount additional training and therefore prohibits overfitting to subjective data sets. There is little training needed for the HMM used by Jump Alignment, but apart from that, neither the tab or MIDI subsystem needs to be trained.

8.1.3 Extended Isophonics data set

As a third contribution, I extended the widely used Isophonics Reference Annotations with 770 MIDI files and 1668 tabs. These symbolic music files are scraped from the Internet and manually matched to the corresponding song in the Isophonics data set. In this research project, we learned that MIDI files and tabs are highly useful for the ACE task. I expect that the extended data set can also contribute to advances in other MIR tasks, such as structural segmentation or lyrics-to-audio alignment.

8.1.4 Python implementations

As a fourth and final contribution, I made my Python implementation of DECIBEL available to the MIR community. The complete code repository can be found on <https://github.com/Daphne0/DECIBEL>. As part of the complete DECIBEL system, this repository contains new Python implementations for existing algorithms, such as Jump Alignment (McVicar et al., 2011b) and data fusion (Koops et al., 2016), as well as implementations for new algorithms, such as my MIDI chord recognizer CASSETTE and a tab parser.

8.2 Future work

From the results of this research project, we can conclude that the integration of multiple symbolic formats and audio for the improvement of ACE is an interesting research direction that deserves more attention in the future. In this section, I give some directions for future work.

As a first suggestion for future work, it would be interesting to test DECIBEL's performance on a larger data set, for example the Billboard data set. The Isophonics data set was convenient to use, because there were myriad MIDI files and tabs available for music by The Beatles and Queen. DECIBEL's performance may drop when testing on less popular or more modern artists, because there may be less symbolic formats available for their songs. It would also be interesting to test DECIBEL on music of another genre, for example jazz music. Very recently, a new data set for jazz, consisting of 113 tracks, was released by Eremenko et al. (2018). In the accompanying paper, Eremenko et al. (2018) evaluate two audio ACE methods on their new data set and observe that the CSR's for these methods were quite low. Extending the data set with MIDI and tab files for these jazz songs and running the DECIBEL system on the extended data set may give better results.

As a second suggestion, I recommend using a larger chord vocabulary in future work. It would be interesting to see if the integration of symbolic formats also helps in recognizing more complex chords, such as seventh chords. Using MIDI and tab files might be very helpful in recognizing complex chords, as these representations directly encode notes and chord labels respectively.

My third suggestion for future work is to experiment various techniques for DECIBEL's subtasks. In the current implementation, I decided to use state-of-the-art audio ACE methods in combination with relatively simple methods for alignment and MIDI chord estimation. The rationale behind this is as follows: the audio ACE methods should be as "good" as possible, as it may be trivial to improve a poorly performing audio ACE method. On the other hand, the techniques for tab alignment, MIDI alignment and MIDI chord recognition should be simple: if we can show that simple methods already improve the audio ACE method, we can assume that more sophisticated techniques will give even better results. Besides, simpler methods are easier to understand and implement, and typically do not require a lot of training. However, as we now know that applying simple

techniques on MIDI and tab files already improves audio ACE, it would be interesting to see how much we can improve our results by using more sophisticated techniques. For example, a MIDI chord recognizer that takes the musical context into account will probably give better results than our proposed CASSETTE algorithm.

As a fourth and final suggestion for future work, I propose to test DECIBEL's performance by doing user tests. We saw that DECIBEL significantly improves each of the seven tested state-of-the-art audio ACE systems. Based on these results, we expect that performers prefer DECIBEL's estimated chord label sequences to the sequences found by audio ACE systems. However, recall from Section 3.3 that WCSR is not always a good measure for the quality of a chord sequence. Therefore, it would be good to verify the improvement by doing user tests.

8.3 Acknowledgements

Having arrived at the *coda* of this thesis, I would like to express my gratitude to my supervisors Anja Volk and Hendrik Vincent Kooops for their inspiring ideas, motivating suggestions and overall great supervision during the development of this thesis, but also in an earlier experimentation project. They have introduced me to the MIR community by inviting me to the Lorentz workshop on Computational Ethnomusicology and thanks to their support I was able to present the results of my experimentation project at the 7th International Workshop on Folk Music Analysis in Málaga. Anja encouraged me to apply for a WiMIR travel grant, which gave me the wonderful opportunity to present my work on DECIBEL as a Late Breaking Demo at the 19th International Society for Music Information Retrieval Conference in Paris. I would also like to thank Matthias Mauch, who gave me a lot of valuable advice during our mentoring sessions in the WiMIR mentoring program. Furthermore, I would like to thank to Frans Wiering for taking on the role of examiner. Thanks to Lucille Mattijssen, who advised me on statistical tests. Finally, I wish to thank my family for their support and encouragement throughout my study.

Appendix A

Data set

<i>The Beatles - Please Please Me</i>	
Index	Song title
1	I Saw Her Standing There
2	Misery
3	Anna (Go To Him)
4	Chains
5	Boys
6	Ask Me Why
7	Please Please Me
8	Love Me Do
9	P. S. I Love You
10	Baby It's You
11	Do You Want To Know A Secret
12	A Taste Of Honey
13	There's A Place
14	Twist And Shout

<i>The Beatles - With the Beatles</i>	
Index	Song title
15	It Won't Be Long
16	All I've Got To Do
17	All My Loving
18	Don't Bother Me
19	Little Child
20	Till There Was You
21	Please Mister Postman
22	Roll Over Beethoven
23	Hold Me Tight
24	You Really Got A Hold On Me
25	I Wanna Be Your Man
26	Devil In Her Heart
27	Not A Second Time
28	Money

The Beatles - A Hard Day's Night

Index	Song title
29	A Hard Day's Night
30	I Should Have Known Better
31	If I Fell
32	I'm Happy Just To Dance With You
33	And I Love Her
34	Tell Me Why
35	Can't Buy Me Love
36	Any Time At All
37	I'll Cry Instead
38	Things We Said Today
39	When I Get Home
40	You Can't Do That
41	I'll Be Back
<i>The Beatles - Beatles for Sale</i>	
Index	Song title
42	No Reply
43	I'm a Loser
44	Baby's In Black
45	Rock and Roll Music
46	I'll Follow the Sun
47	Mr. Moonlight
48	Kansas City- Hey, Hey, Hey, Hey
49	Eight Days a Week
50	Words of Love
51	Honey Don't
52	Every Little Thing
53	I Don't Want to Spoil the Party
54	What You're Doing
55	Everybody's Trying to Be My Baby
<i>The Beatles - Help!</i>	
Index	Song title
56	Help!
57	The Night Before
58	You've Got To Hide Your Love Away
59	I Need You
60	Another Girl
61	You're Going To Lose That Girl
62	Ticket To Ride
63	Act Naturally
64	It's Only Love
65	You Like Me Too Much
66	Tell Me What You See
67	I've Just Seen a Face
68	Yesterday
69	Dizzy Miss Lizzy
<i>The Beatles - Rubber Soul</i>	

Index	Song title
70	Drive My Car
71	Norwegian Wood (This Bird Has Flown)
72	You Won't See Me
73	Nowhere Man
74	Think For Yourself
75	The Word
76	Michelle
77	What Goes On
78	Girl
79	I'm Looking Through You
80	In My Life
81	Wait
82	If I Needed Someone
83	Run For Your Life
<i>The Beatles - Revolver</i>	
Index	Song title
84	Taxman
85	Eleanor Rigby
86	I'm Only Sleeping
87	Love You To
88	Here, There And Everywhere
89	Yellow Submarine
90	She Said She Said
91	Good Day Sunshine
92	And Your Bird Can Sing
93	For No One
94	Doctor Robert
95	I Want To Tell You
96	Got To Get You Into My Life
97	Tomorrow Never Knows
<i>The Beatles - Sgt. Pepper's Lonely Hearts Club Band</i>	
Index	Song title
98	Sgt. Pepper's Lonely Hearts Club Band
99	With A Little Help From My Friends
100	Lucy In The Sky With Diamonds
101	Getting Better
102	Fixing A Hole
103	She's Leaving Home
104	Being For The Benefit Of Mr. Kite!
105	Within You Without You
106	When I'm Sixty-Four
107	Lovely Rita
108	Good Morning Good Morning
109	Sgt. Pepper's Lonely Hearts Club Band (Reprise)
110	A Day In The Life
<i>The Beatles - Magical Mystery Tour</i>	

Index	Song title
111	Magical Mystery Tour
112	The Fool On The Hill
113	Flying
114	Blue Jay Way
115	Your Mother Should Know
116	I Am The Walrus
117	Hello Goodbye
118	Strawberry Fields Forever
119	Penny Lane
120	Baby You're A Rich Man
121	All You Need Is Love
<i>The Beatles - The Beatles</i>	
Index	Song title
122	Back in the USSR
123	Dear Prudence
124	Glass Onion
125	Ob-La-Di, Ob-La-Da
126	Wild Honey Pie
127	The Continuing Story of Bungalow Bill
128	While My Guitar Gently Weeps
129	Happiness is a Warm Gun
130	Martha My Dear
131	I'm So Tired
132	Black Bird
133	Piggies
134	Rocky Raccoon
135	Don't Pass Me By
136	Why Don't We Do It In The Road
137	I Will
138	Julia
139	Birthday
140	Yer Blues
141	Mother Nature's Son
142	Everybody's Got Something To Hide Except Me and My Monkey
143	Sexy Sadie
144	Helter Skelter
145	Long Long Long
146	Revolution
147	Honey Pie
148	Savoy Truffle
149	Cry Baby Cry
150	Revolution
151	Good Night
<i>The Beatles - Abbey Road</i>	
Index	Song title
152	Come Together

153	Something
154	Maxwell's Silver Hammer
155	Oh! Darling
156	Octopus's Garden
157	I Want You
158	Here Comes The Sun
159	Because
160	You Never Give Me Your Money
161	Sun King
162	Mean Mr Mustard
163	Polythene Pam
164	She Came In Through The Bathroom Window
165	Golden Slumbers
166	Carry That Weight
167	The End
168	Her Majesty

The Beatles - Let It Be

Index	Song title
169	Two of Us
170	Dig a Pony
171	Across the Universe
172	I Me Mine
173	Dig It
174	Let It Be
175	Maggie Mae
176	I've Got A Feeling
177	One After
178	The Long and Winding Road
179	For You Blue
180	Get Back

Queen - Greatest Hits I

Index	Song title
181	Bohemian Rhapsody
182	Another One Bites The Dust
184	Fat Bottomed Girls
185	Bicycle Race
186	You're My Best Friend
187	Don't Stop Me Now
188	Save Me
189	Crazy Little Thing Called Love
190	Somebody To Love
192	Good Old Fashioned Lover Boy
193	Play The Game
195	Seven Seas Of Rhye
196	We Will Rock You
197	We Are The Champions

Queen - Greatest Hits II

Index	Song title
198	A Kind Of Magic
201	I Want It All
202	I Want To Break Free
206	Who Wants To Live Forever
211	Hammer To Fall
212	Friends Will Be Friends

Appendix B

Alignment listening test results

MIDI file	Alignment error	Listening test result
128-005	0.723	2 - Good alignment, no fade out in MIDI
014-003	0.762	2 - Good alignment, MIDI cut off too soon at the end
095-002	0.636	2 - Good alignment, but seems to be a bit shifted
047-001	0.934	1 - Very bad alignment
063-004	0.764	3 - Good alignment
098-004	0.850	3 - Good alignment
196-001	0.656	3 - Good alignment
099-002	0.959	1 - Very bad alignment
195-001	0.723	2 - Good alignment with minor issues at the end
036-002	0.903	1 - First verse and chorus missing in MIDI, issues at the end
184-001	0.720	3 - Good alignment
078-003	0.654	2 - Good alignment, no fade out in MIDI
042-001	0.815	2 - Good alignment, but sometimes seems to be a bit shifted
152-002	0.705	2 - Good alignment, no fade out in MIDI
181-005	0.964	1 - Very bad alignment (transcription with only piano)
002-001	0.644	2 - Good alignment, no fade out in MIDI
061-002	0.762	2 - Good alignment, minor issues at the end
155-001	0.880	1 - Mediocre alignment, last 50s of the audio missing in MIDI file
073-004	0.760	2 - Good alignment, but seems to be a bit shifted
055-002	0.712	3 - Good alignment
003-001	0.771	2 - Good alignment, but seems to be a bit shifted
010-001	0.709	2 - Good alignment, no fade out in MIDI
125-002	0.784	3 - Good alignment
202-005	0.641	3 - Good alignment
212-001	0.532	3 - Good alignment

Table B.1: Results of the alignment listening test

Appendix C

Tabs in parsing evaluation set

The annotated subset for the tab parsing evaluation consists of these 25 songs:

1. Let It Be (ver 6) Tabs
2. A Hard Days Night (ver 10) Chords
3. You're My Best Friend solo Tabs
4. Savoy Truffle (ver 2) Chords
5. Strawberry Fields Forever Acoustic Chords
6. A Taste Of Honey (ver 2) Chords
7. What You're Doing Chords
8. A Hard Days Night (ver 3) Tabs
9. Let It Be (ver 4) Chords
10. Love You To Tabs
11. Cry Baby Cry (ver 4) Chords
12. For You Blue Chords
13. Getting Better (ver 5) Chords
14. Help (ver 8) Chords
15. No Reply (ver 4) Chords
16. Happiness Is A Warm Gun Tabs
17. I'll Be Back (ver 2) Tabs
18. Tell Me Why Tabs

19. Get Back (ver 5) Tabs
20. Anna (ver 2) Chords
21. A Day In The Life Tabs
22. You Never Give Me Your Money Tabs
23. Eleanor Rigby (ver 3) Tabs
24. I'll Be Back (ver 7) Chords
25. Yesterday Chords

Appendix D

Tabs in forward and backward training set

The following 22 songs were used for training forward and backward probabilities:

1. Misery Chords
2. Please Mister Postman Chords
3. You Really Got A Hold On Me (ver 2) Chords
4. And I Love Her Tabs
5. I'll Cry Instead Tabs
6. I'll Be Back Tabs
7. Mr Moonlight Tabs
8. Help (ver 6) Chords
9. I Need You Tabs
10. You're Going To Lose That Girl Chords
11. Act Naturally (ver 3) Tabs
12. Norwegian Wood (ver 4) Chords
13. Think For Yourself Chords
14. Girl Acoustic Chords
15. Good Day Sunshine Chords
16. Sgt Peppers Lonely Hearts Club Band Reprise (ver 2) Chords
17. Penny Lane (ver 6) Chords

18. Come Together Chords
19. Mean Mr Mustard (ver 2) Chords
20. Golden Slumbers Chords
21. A Kind Of Magic Chords
22. I Want To Break Free Chords

Appendix E

Tukey’s HSD test results

E.1 CHF

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	4.46e-07	9.98e-08				
MV-ALL	0.102	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.000225		
DF-ALL	1	3.71e-08	2.59e-07	0.132	3.71e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	1.09e-07	0.738	3.71e-08

Table E.1: P-values Tukey’s HSD test; Audio ACE: CHF, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	9.48e-05				
MV-ALL	0.000378	3.71e-08	3.71e-08			
MV-BEST	0.709	3.71e-08	3.71e-08	0.0932		
DF-ALL	3.71e-08	3.71e-08	3.71e-08	0.00054	3.71e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	3.08e-07	3.71e-08	0.752

Table E.2: P-values Tukey’s HSD test; Audio ACE: CHF, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000845				
MV-ALL	0.000132	3.71e-08	3.71e-08			
MV-BEST	2.43e-07	3.71e-08	4.39e-08	3.71e-08		
DF-ALL	3.71e-08	3.71e-08	3.71e-08	0.0001	3.71e-08	
DF-BEST	0.891	3.71e-08	3.71e-08	0.0167	3.71e-08	3.71e-08

Table E.3: P-values Tukey’s HSD test; Audio ACE: CHF, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	4.63e-08				
MV-ALL	0.364	3.71e-08	3.71e-08			
MV-BEST	0.0663	3.71e-08	3.71e-08	1.71e-05		
DF-ALL	5.43e-05	3.71e-08	3.75e-08	0.121	3.71e-08	
DF-BEST	0.00533	3.71e-08	3.71e-08	2.73e-07	0.987	3.71e-08

Table E.4: P-values Tukey's HSD test; Audio ACE: CHF, Measure: Segmentation

E.2 CM2

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	4.25e-06					
RND-BEST	1	1.93e-05				
MV-ALL	3.71e-08	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.891		
DF-ALL	3.71e-08	3.71e-08	3.71e-08	0.096	0.00173	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	0.641	0.999	0.00025

Table E.5: P-values Tukey's HSD test; Audio ACE: CM2, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000118				
MV-ALL	0.998	3.71e-08	3.71e-08			
MV-BEST	0.31	3.71e-08	3.71e-08	0.0878		
DF-ALL	1.01e-06	3.71e-08	3.71e-08	2.05e-05	3.71e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	3.71e-08	3.71e-08	0.118

Table E.6: P-values Tukey's HSD test; Audio ACE: CM2, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000845				
MV-ALL	0.00393	3.71e-08	3.71e-08			
MV-BEST	8.88e-06	3.71e-08	3.72e-08	3.71e-08		
DF-ALL	3.71e-08	3.71e-08	3.71e-08	1.13e-05	3.71e-08	
DF-BEST	0.201	3.71e-08	3.71e-08	0.846	3.71e-08	4.21e-08

Table E.7: P-values Tukey's HSD test; Audio ACE: CM2, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	1.2e-07				
MV-ALL	0.132	3.71e-08	3.71e-08			
MV-BEST	0.799	3.71e-08	3.71e-08	0.00131		
DF-ALL	6.98e-06	3.71e-08	3.71e-08	0.151	3.85e-08	
DF-BEST	0.724	3.71e-08	3.71e-08	0.000767	1	3.77e-08

Table E.8: P-values Tukey's HSD test; Audio ACE: CM2, Measure: Segmentation

E.3 JLW1

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	0.000359	1.46e-06				
MV-ALL	0.000596	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.00199		
DF-ALL	0.586	3.71e-08	6.96e-08	0.186	4.25e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	2.16e-07	0.499	3.71e-08

Table E.9: P-values Tukey's HSD test; Audio ACE: JLW1, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000146				
MV-ALL	3.71e-08	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	1		
DF-ALL	0.0959	3.71e-08	3.71e-08	8.95e-05	0.000172	
DF-BEST	0.0662	3.71e-08	3.71e-08	3.71e-08	3.71e-08	6.04e-07

Table E.10: P-values Tukey's HSD test; Audio ACE: JLW1, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.00113				
MV-ALL	1	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.72e-08	3.71e-08		
DF-ALL	2.29e-05	3.71e-08	3.71e-08	8.84e-06	3.71e-08	
DF-BEST	1	3.71e-08	3.71e-08	1	3.71e-08	6.14e-06

Table E.11: P-values Tukey's HSD test; Audio ACE: JLW1, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	3.71e-08				
MV-ALL	0.378	3.71e-08	3.71e-08			
MV-BEST	0.0494	3.71e-08	3.71e-08	1.13e-05		
DF-ALL	2.16e-05	3.71e-08	3.47e-07	0.0706	3.71e-08	
DF-BEST	0.97	3.71e-08	3.71e-08	0.0462	0.393	1.73e-07

Table E.12: P-values Tukey's HSD test; Audio ACE: JLW1, Measure: Segmentation

E.4 JLW2

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	0.0108	3.54e-06				
MV-ALL	5.72e-07	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.00411		
DF-ALL	0.0447	3.71e-08	3.98e-08	0.132	4.4e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	5.81e-06	0.779	3.71e-08

Table E.13: P-values Tukey's HSD test; Audio ACE: JLW2, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000224				
MV-ALL	3.71e-08	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	1		
DF-ALL	0.0337	3.71e-08	3.71e-08	7.57e-05	4.07e-05	
DF-BEST	0.211	3.71e-08	3.71e-08	3.71e-08	3.71e-08	1.06e-06

Table E.14: P-values Tukey's HSD test; Audio ACE: JLW2, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000277				
MV-ALL	0.26	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.74e-08	3.72e-08		
DF-ALL	0.108	3.71e-08	3.71e-08	1.71e-05	3.71e-08	
DF-BEST	0.61	3.71e-08	3.71e-08	0.998	3.71e-08	0.000249

Table E.15: P-values Tukey's HSD test; Audio ACE: JLW2, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	3.72e-08				
MV-ALL	1	3.71e-08	3.71e-08			
MV-BEST	0.000202	3.71e-08	3.71e-08	0.000124		
DF-ALL	0.0362	3.71e-08	4.5e-08	0.0493	3.71e-08	
DF-BEST	0.357	3.71e-08	3.71e-08	0.297	0.232	5.43e-06

Table E.16: P-values Tukey's HSD test; Audio ACE: JLW2, Measure: Segmentation

E.5 KBK1

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	8.4e-08				
MV-ALL	1	3.71e-08	3.71e-08			
MV-BEST	7.4e-06	3.71e-08	3.71e-08	5.44e-05		
DF-ALL	0.35	3.71e-08	6.16e-06	0.159	3.71e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	3.76e-08	0.445	3.71e-08

Table E.17: P-values Tukey's HSD test; Audio ACE: KBK1, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000106				
MV-ALL	3.71e-08	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.877		
DF-ALL	0.000842	3.71e-08	3.71e-08	0.000111	0.0167	
DF-BEST	0.731	3.71e-08	3.71e-08	3.71e-08	3.71e-08	4.42e-07

Table E.18: P-values Tukey's HSD test; Audio ACE: KBK1, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.00022				
MV-ALL	0.399	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.97e-08	3.71e-08		
DF-ALL	0.054	3.71e-08	3.71e-08	1.57e-05	3.71e-08	
DF-BEST	0.15	3.71e-08	3.71e-08	0.999	3.71e-08	1.1e-06

Table E.19: P-values Tukey's HSD test; Audio ACE: KBK1, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	3.71e-08				
MV-ALL	1	3.71e-08	3.71e-08			
MV-BEST	1.83e-06	3.71e-08	3.71e-08	3.87e-06		
DF-ALL	0.0442	3.71e-08	2.09e-07	0.029	3.71e-08	
DF-BEST	0.0216	3.71e-08	3.71e-08	0.0334	0.335	4.74e-08

Table E.20: P-values Tukey's HSD test; Audio ACE: KBK1, Measure: Segmentation

E.6 KBK2

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	3.71e-08				
MV-ALL	3.22e-06	3.71e-08	1.39e-05			
MV-BEST	0.577	3.71e-08	3.71e-08	3.71e-08		
DF-ALL	3.71e-08	3.71e-08	0.0871	0.283	3.71e-08	
DF-BEST	4.23e-05	3.71e-08	3.71e-08	3.71e-08	0.0442	3.71e-08

Table E.21: P-values Tukey's HSD test; Audio ACE: KBK2, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000117				
MV-ALL	3.71e-08	3.71e-08	3.71e-08			
MV-BEST	5.64e-08	3.71e-08	3.71e-08	0.641		
DF-ALL	0.0209	3.71e-08	3.71e-08	0.000111	0.06	
DF-BEST	0.0127	3.71e-08	3.71e-08	3.71e-08	3.71e-08	3.78e-08

Table E.22: P-values Tukey's HSD test; Audio ACE: KBK2, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.00102				
MV-ALL	0.343	3.71e-08	3.71e-08			
MV-BEST	3.76e-08	3.71e-08	3.73e-08	3.71e-08		
DF-ALL	3.71e-08	3.71e-08	3.71e-08	5.39e-06	3.71e-08	
DF-BEST	1	3.71e-08	3.71e-08	0.385	3.74e-08	3.71e-08

Table E.23: P-values Tukey's HSD test; Audio ACE: KBK2, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	3.73e-08				
MV-ALL	5.1e-05	3.71e-08	3.71e-08			
MV-BEST	0.35	3.71e-08	3.71e-08	3.72e-08		
DF-ALL	3.71e-08	3.71e-08	0.000171	0.0362	3.71e-08	
DF-BEST	0.993	3.71e-08	3.71e-08	1.37e-06	0.798	3.71e-08

Table E.24: P-values Tukey's HSD test; Audio ACE: KBK2, Measure: Segmentation

E.7 WL1

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	1.13e-05	6.91e-07				
MV-ALL	0.0304	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.000342		
DF-ALL	0.999	3.71e-08	8.89e-07	0.115	3.71e-08	
DF-BEST	3.71e-08	3.71e-08	3.71e-08	8.1e-08	0.626	3.71e-08

Table E.25: P-values Tukey's HSD test; Audio ACE: WL1, Measure: Chord Symbol Recall

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000182				
MV-ALL	3.71e-08	3.71e-08	3.71e-08			
MV-BEST	3.74e-08	3.71e-08	3.71e-08	1		
DF-ALL	0.0546	3.71e-08	3.71e-08	0.000489	0.00251	
DF-BEST	0.0905	3.71e-08	3.71e-08	3.71e-08	3.71e-08	3.7e-07

Table E.26: P-values Tukey's HSD test; Audio ACE: WL1, Measure: Oversegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	3.71e-08	0.000202				
MV-ALL	3.72e-08	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.96e-08	1.14e-07		
DF-ALL	0.634	3.71e-08	3.71e-08	5.45e-06	3.71e-08	
DF-BEST	1.38e-06	3.71e-08	3.71e-08	0.792	3.71e-08	0.0036

Table E.27: P-values Tukey's HSD test; Audio ACE: WL1, Measure: Undersegmentation

	Original	RND-ALL	RND-BEST	MV-ALL	MV-BEST	DF-ALL
RND-ALL	3.71e-08					
RND-BEST	1.22e-06	3.71e-08				
MV-ALL	6.09e-05	3.71e-08	3.71e-08			
MV-BEST	3.71e-08	3.71e-08	3.71e-08	0.000182		
DF-ALL	0.717	3.71e-08	3.71e-08	0.0293	3.71e-08	
DF-BEST	8.76e-08	3.71e-08	3.71e-08	0.857	0.0273	0.000202

Table E.28: P-values Tukey's HSD test; Audio ACE: WL1, Measure: Segmentation

Bibliography

- Ritesh Ajoodha, Richard Klein, and Benjamin Rosman. Single-labelled music genre classification using content-based features. In *Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech), 2015*, pages 66–71. IEEE, 2015.
- Andreas Arzt. *Flexible and Robust Music Tracking*. PhD thesis, Johannes Kepler University Linz, 2016.
- Juan Pablo Bello and Jeremy Pickens. A robust mid-level representation for harmonic content in music signals. In *Proceedings of the 6th International Society for Music Information Retrieval Conference*, pages 304–311, 2005.
- Alessandro Bonvivi. Automatic chord recognition using deep learning techniques. Master’s thesis, Politecnico di Milano, 2014.
- Dimitrios Bountouridis, Hendrik Vincent Koops, Frans Wiering, and Remco C. Veltkamp. A data-driven approach to chord similarity and chord mutability. In *Second International Conference on Multimedia Big Data*, pages 275–278. IEEE, 2016.
- Judith C. Brown. Calculation of a constant q spectral transform. *The Journal of the Acoustical Society of America*, 89(1):425–434, 1991.
- John Ashley Burgoyne, Laurent Pugin, I. Fujinaga, and Kereliuk Ichiro. A cross-validated study of modelling strategies for automatic chord recognition in audio. In *Proceedings of the 8th International Society for Music Information Retrieval Conference*, pages 251–254. ISMIR, 2007.
- Chris Cannam, Emmanouil Benetos, Matthias Mauch, Matthew E.P. Davies, Simon Dixon, Christian Landone, Katy Noland, and Dan Stowell. MIREX 2018: Vamp plugins from the centre for digital music. *Proceedings of the Music Information Retrieval Evaluation eXchange (MIREX)*, 2017.
- Julio José Carabias-Orti, Francisco J. Rodríguez-Serrano, Pedro Vera-Candeas, Nicolás Ruiz-Reyes, and Francisco J. Cañadas-Quesada. An audio to score alignment framework using spectral factorization and dynamic time warping. In *Proceedings of the 16th International Society for Music Information Retrieval Conference*, pages 742–748. ISMIR, 2015.
- Michael Scott Cuthbert and Christopher Ariza. music21: A toolkit for computer-aided musicology and symbolic music data. In *Proceedings of the 11th International Society for Music Information Retrieval Conference*, pages 637–642. ISMIR, 2010.

- Philippe Cuvillier and Arshia Cont. Coherent time modeling of semi-markov models with application to real-time audio-to-score alignment. In *Machine Learning for Signal Processing (MLSP), 2014 IEEE International Workshop on*, pages 1–6. IEEE, 2014.
- Simon Dixon and Gerhard Widmer. Match: A music alignment tool chest. In *Proceedings of the 6th International Society for Music Information Retrieval Conference*, pages 492–497, 2005.
- Tuomas Eerola and Petri Toiviainen. MIR in Matlab: The MIDI toolbox. In *Proceedings of the 5th International Society for Music Information Retrieval Conference*, pages 22–27. ISMIR, 2004.
- Daniel P.W. Ellis and Graham E. Poliner. Identifying cover songs’ with chroma features and dynamic programming beat tracking. In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2007*, volume 4, pages 1429–1432. IEEE, 2007.
- Vsevolod Eremenko, Emir Demirel, Baris Bozkurt, and Xavier Serra. Audio-aligned jazz harmony dataset for automatic chord transcription and corpus-based research. In *Proceedings of the 19th International Society for Music Information Retrieval Conference*, pages 483–490. ISMIR, 2018.
- Sebastian Ewert, Meinard Müller, Verena Konz, Daniel Müllensiefen, and Geraint A. Wiggins. Towards cross-version harmonic analysis of music. *IEEE Transactions on Multimedia*, 14(3):770–782, 2012.
- Milton Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the american statistical association*, 32(200):675–701, 1937.
- Takuya Fujishima. Realtime chord recognition of musical sound: a system using Common Lisp Music. In *Proceedings of the International Computer Music Conference (ICMC)*, pages 464–467, 1999.
- Dennis Gabor. Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering*, 93(26):429–441, 1946.
- Emilia Gómez and Perfecto Herrera. The song remains the same: identifying versions of the same piece using tonal descriptors. In *Proceedings of the 7th International Society for Music Information Retrieval Conference*, pages 180–185, 2006.
- Robert Guérin. *MIDI power!: The comprehensive guide*. Course Technology PTR, 2009.
- Christopher Harte. *Towards automatic extraction of harmony information from music signals*. PhD thesis, Department of Electronic Engineering, Queen Mary, University of London, 2010.
- Christopher Harte and Mark Sandler. Automatic chord identification using a quantised chromagram. In *Proceedings of the 118th Audio Engineering Society Convention*, pages 745–750. AES, 2005.

- Christopher Harte, Mark B. Sandler, Samer A. Abdallah, and Emilia Gómez. Symbolic representation of musical chords: A proposed syntax for text annotations. In *Proceedings of the 6th International Society for Music Information Retrieval Conference*, pages 66–71. ISMIR, 2005.
- Christopher Harte, Mark Sandler, and Martin Gasser. Detecting harmonic change in musical audio. In *Proceedings of the 1st ACM Workshop on Audio and Music Computing Multimedia*, pages 21–26. ACM, 2006.
- Jerry L. Hintze and Ray D. Nelson. Violin plots: a box plot-density trace synergism. *The American Statistician*, 52(2):181–184, 1998.
- Ning Hu, Roger B. Dannenberg, and George Tzanetakis. Polyphonic audio matching and alignment for music retrieval. In *Applications of Signal Processing to Audio and Acoustics, 2003 IEEE Workshop on.*, pages 185–188. IEEE, 2003.
- Eric J. Humphrey and Juan Pablo Bello. Four timely insights on automatic chord estimation. In *Proceedings of the 16th International Society for Music Information Retrieval Conference*, pages 673–679, 2015.
- Fumitada Itakura. Minimum prediction residual principle applied to speech recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 23(1):67–72, 1975.
- Junyan Jiang, Wei Li, and Yiming Wu. Extended abstract for MIREX 2017 submission: Chord recognition using random forest model. *Proceedings of the Music Information Retrieval Evaluation eXchange (MIREX)*, 2017.
- Artemy Kolchinsky, Nakul Dhande, Kengjeun Park, and Yong-Yeol Ahn. The minor fall, the major lift: inferring emotional valence of musical chords through lyrics. *Royal Society Open Science*, 4(11), 2017.
- Verena Konz, Meinard Müller, and Rainer Kleinertz. A cross-version chord labelling approach for exploring harmonic structures: A case study on Beethoven’s *Appassionata*. *Journal of New Music Research*, 42(1):61–77, 2013.
- Hendrik Vincent Koops, W. Bas de Haas, Dimitrios Bountouridis, and Anja Volk. Integration and quality assessment of heterogeneous chord sequences using data fusion. In *Proceedings of the 17th International Society for Music Information Retrieval Conference*, pages 178–184. ISMIR, 2016.
- Hendrik Vincent Koops, W. Bas de Haas, Jeroen Bransen, and Anja Volk. Chord label personalization through deep learning of integrated harmonic interval-based representations. In *Proceedings of the First International Workshop on Deep Learning for Music*, pages 19–25, 2017.
- Filip Korzeniowski and Gerhard Widmer. Feature learning for chord recognition: The Deep Chroma Extractor. In *Proceedings of 17th International Conference on Music Information Retrieval*, 2016a.
- Filip Korzeniowski and Gerhard Widmer. A fully convolutional deep auditory model for musical chord recognition. In *Proceedings of the 26th International Workshop on Machine Learning for Signal Processing*, pages 1–6. IEEE, 2016b.

- Rémi Lajugie, Piotr Bojanowski, Philippe Cuvillier, Sylvain Arlot, and Francis Bach. A weakly-supervised discriminative model for audio-to-score alignment. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 2484–2488. IEEE, 2016.
- Mark Lewisohn. *Beatles Recording Sessions: The Official Abbey Road Studio Session Notes, 1962-1970*. Harmony Books, 1989.
- Robert Macrae. *Linking Music Metadata*. PhD thesis, Queen Mary University of London, 2012.
- Robert Macrae and Simon Dixon. Guitar tab mining, analysis and ranking. In *Proceedings of the 12th International Society for Music Information Retrieval Conference*, pages 453–458. ISMIR, 2011.
- Ladislav Maršík, Martin Rusek, Kateřina Slaninová, Jan Martinovič, and Jaroslav Pokorný. Evaluation of chord and chroma features and dynamic time warping scores on cover song identification task. In *IFIP International Conference on Computer Information Systems and Industrial Management*, pages 205–217. Springer, 2017.
- Kristen Masada and Razvan Bunescu. Chord recognition in symbolic music using semi-Markov conditional random fields. In *Proceedings of the 18th International Society for Music Information Retrieval Conference*, pages 272–278. ISMIR, 2017.
- Matthias Mauch. *Automatic Chord Transcription from Audio Using Computational Models of Musical Context*. PhD thesis, Queen Mary, University of London, 2010.
- Matthias Mauch and Simon Dixon. A discrete mixture model for chord labeling. In *Proceedings of the 9th International Society for Music Information Retrieval Conference*, pages 45–50. ISMIR, 2008.
- Matthias Mauch and Simon Dixon. Simultaneous estimation of chords and musical context from audio. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1280–1289, 2010.
- Matthias Mauch, Chris Cannam, Matthew Davies, Simon Dixon, Christopher Harte, Sefki Kolozali, Dan Tidhar, and Mark Sandler. Omras2 metadata project 2009. In *Proceedings of 10th International Conference on Music Information Retrieval*, 2009.
- Matthias Mauch, Hiromasa Fujihara, and Masataka Goto. Integrating additional chord information into HMM-based lyrics-to-audio alignment. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(1):200–210, 2012.
- H. John Maxwell. An expert system for harmonic analysis of tonal music. In *Proceedings of the First Workshop on AI and Music*, pages 20–33. AAAI, 1992.
- Brian McFee and Juan Pablo Bello. Structured training for large-vocabulary chord recognition. In *Proceedings of the 18th International Society for Music Information Retrieval Conference*, 2017.

- Brian McFee, Matt McVicar, Stefan Balke, Carl Thomé, Colin Raffel, Dana Lee, Oriol Nieto, Eric Battenberg, Dan Ellis, Ryuichi Yamamoto, Josh Moore, Rachel Bittner, Keunwoo Choi, Pius Friesch, Fabian-Robert Stöter, Vincent Lostanlen, Siddhartha Kumar, Simon Waloschek, Seth, Rimvydas Naktinis, Douglas Repetto, Curtis "Fjord" Hawthorne, CJ Carr, Waldir Pimenta, Petr Viktorin, Paul Brossier, João Felipe Santos, JackieWu, Erik, and Adrian Holovaty. *librosa/librosa*: 0.6.1, May 2018. URL <https://doi.org/10.5281/zenodo.1252297>.
- Cory McKay and Ichiro Fujinaga. *jSymbolic*: A feature extractor for MIDI files. In *Proceedings of the 2006 International Computer Music Conference*, pages 302–305. ICMC, 2006.
- Matt McVicar and Tijl De Bie. Enhancing chord recognition accuracy using web resources. In *Proceedings of the 3rd International Workshop on Machine Learning and Music*, pages 41–44. ACM, 2010.
- Matt McVicar, Yizhao Ni, Tijl De Bie, and Raul Santos-Rodriguez. Leveraging noisy online databases for use in chord recognition. In *Proceedings of the 12th International Society for Music Information Retrieval Conference*, pages 639–644. ISMIR, 2011a.
- Matt McVicar, Yizhao Ni, Raul Santos-Rodriguez, and Tijl De Bie. Using online chord databases to enhance chord recognition. *Journal of New Music Research*, 40(2):139–152, 2011b.
- Matt McVicar, Raúl Santos-Rodríguez, Yizhao Ni, and Tijl De Bie. Automatic chord estimation from audio: A review of the state of the art. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 22(2):556–575, 2014.
- Meinard Müller. *Fundamentals of Music Processing: Audio, Analysis, Algorithms, Applications*. Springer International Publishing, 2015.
- Meinard Müller and Sebastian Ewert. Towards timbre-invariant audio features for harmony-based music. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(3):649–662, 2010.
- Yizhao Ni, Matt McVicar, Raul Santos-Rodriguez, and Tijl De Bie. An end-to-end machine learning system for harmonic analysis of music. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(6):1771–1783, 2012.
- Yizhao Ni, Matt McVicar, Raul Santos-Rodriguez, and Tijl De Bie. Understanding effects of subjectivity in measuring chord estimation accuracy. *IEEE Transactions on Audio, Speech, and Language Processing*, 21(12):2607–2615, 2013.
- Ken O’Hanlon, Sebastian Ewert, Johan Pauwels, and Mark B. Sandler. Improved template based chord recognition using the crp feature. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 306–310. IEEE, 2017.
- Hélele Papadopoulos and George Tzanetakis. Modeling chord and key structure with markov logic. In *Proceedings of the 13th International Society for Music Information Retrieval Conference*, pages 127–132, 2012.

- Brian Pardo and William P. Birmingham. The chordal analysis of tonal music. Technical Report CSE-TR-439-01, The University of Michigan, Department of Electrical Engineering and Computer Science, 2001.
- Bryan Pardo and William P. Birmingham. Algorithms for chordal analysis. *Computer Music Journal*, 26(2):27–49, 2002.
- Steffen Pauws. Musical key extraction from audio. In *Proceedings of the 5th International Society for Music Information Retrieval Conference*, pages 96–99. ISMIR, 2004.
- Thomas Prätzlich and Meinard Müller. Triple-based analysis of music alignments without the need of ground-truth annotations. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 266–270. IEEE, 2016.
- Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- Daniele Radicioni and Roberto Esposito. BREVE: an HMPerceptron-based chord recognition system. *Advances in Music Information Retrieval*, pages 143–164, 2010.
- Colin Raffel. *Learning-Based Methods for Comparing Sequences, with Applications to Audio-to-MIDI Alignment and Matching*. PhD thesis, Columbia University, 2016.
- Colin Raffel and Daniel P.W. Ellis. Intuitive analysis, creation and manipulation of MIDI data with pretty_midi. In *15th International Society for Music Information Retrieval Conference Late Breaking and Demo Papers*, pages 84–93. ISMIR, 2014.
- Colin Raffel and Daniel P.W. Ellis. Large-scale content-based matching of midi and audio files. In *Proceedings of the 16th International Society for Music Information Retrieval Conference*, pages 234–240. ISMIR, 2015.
- Colin Raffel and Daniel P.W. Ellis. Extracting ground-truth information from MIDI files: A MIDifesto. In *Proceedings of the 17th International Society for Music Information Retrieval Conference*, pages 796–802. ISMIR, 2016a.
- Colin Raffel and Daniel P.W. Ellis. Optimizing dtw-based audio-to-midi alignment and matching. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 81–85. IEEE, 2016b.
- Christopher Raphael and Joshua Stoddard. Harmonic analysis with probabilistic graphical models. In *Proceedings of the 4th International Society for Music Information Retrieval Conference*, pages 177–181. ISMIR, 2003.
- Christopher Raphael and Joshua Stoddard. Functional harmonic analysis using probabilistic models. *Computer Music Journal*, 28(3):45–52, 2004.
- Jeremy Reed, Yushi Ueda, Sabato Marco Siniscalchi, Yuuki Uchiyama, Shigeki Sagayama, and Chin-Hui Lee. Minimum classification error training to improve isolated chord recognition. In *Proceedings of the 10th International Society for Music Information Retrieval Conference*, pages 609–614. ISMIR, 2009.

- Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing*, 26(1): 43–49, 1978.
- Ricardo Scholz and Geber Ramalho. COCHONUT: Recognizing complex chords from MIDI guitar sequences. In *Proceedings of the 9th International Society for Music Information Retrieval Conference*, pages 27–32. ISMIR, 2008.
- Ricardo Scholz, Emmanuel Vincent, and Frédéric Bimbot. Robust modeling of musical chord sequences using probabilistic n-grams. In *Proceedings of the 34th International Conference on Acoustics, Speech and Signal Processing*, pages 53–56. IEEE, 2009.
- Ricardo Scholz, Geber Ramalho, and Giordano Cabral. Cross task study on MIREX recent results: An index for evolution measurement and some stagnation hypotheses. In *Proceedings of the 17th International Society for Music Information Retrieval Conference*, pages 372–378, 2016.
- Alexander Sheh and Daniel P.W. Ellis. Chord segmentation and recognition using em-trained hidden Markov models. In *Proceedings of the 4th International Society for Music Information Retrieval Conference*, pages 183–189. ISMIR, 2003.
- Arun Shenoy and Ye Wang. Key, chord, and rhythm tracking of popular music recordings. *Computer Music Journal*, 29(3):75–86, 2005.
- Siddharth Sigtia, Nicolas Boulanger-Lewandowski, and Simon Dixon. Audio chord recognition with a hybrid recurrent neural network. In *Proceedings of the 16th International Society for Music Information Retrieval Conference*, pages 127–133. ISMIR, 2015.
- Daniel Sleator and David Temperley. The Melisma music analyzer. *Web: <http://www.link.cs.cmu.edu/music-analysis>*, 2001.
- Dan Stowell and Simon Dixon. MIR in school? Lessons from ethnographic observation of secondary school music classes. In *Proceedings of the 12th International Society for Music Information Retrieval Conference*, pages 347–352. ISMIR, 2011.
- Eric Taylor. *The AB Guide to Music Theory, Part I*. The Associated Board of the Royal Schools of Music (Publishing) Ltd, 1989.
- David Temperley and Daniel Sleator. Modeling meter and harmony: A preference-rule approach. *Computer Music Journal*, 23(1):10–27, 1999.
- John W. Tukey. Comparing individual means in the analysis of variance. *Biometrics*, pages 99–114, 1949.
- Robert J. Turetsky and Daniel P.W. Ellis. Ground-truth transcriptions of real music from force-aligned midi syntheses. In *Proceedings of the 4th International Society for Music Information Retrieval Conference*, pages 135–141, 2003.
- Gregory H. Wakefield. Mathematical representation of joint time-chroma distributions. In *Proceedings of SPIE's International Symposium on Optical Science, Engineering, and Instrumentation*, volume 3807, pages 637–646. International Society for Optics and Photonics, 1999.

Siyang Wang, Sebastian Ewert, and Simon Dixon. Robust joint alignment of multiple versions of a piece of music. In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, page 83. ISMIR, 2014.

Terry Winograd. Linguistics and the computer analysis of tonal harmony. *Journal of Music Theory*, 12(1):2–49, 1968.

Yiming Wu, Xiangyi Feng, and Wei Li. MIREX 2017 submission: Automatic audio chord recognition with MIDI-trained deep feature and BLSTM-CRF sequence decoding model. *Proceedings of the Music Information Retrieval Evaluation eXchange (MIREX)*, 2017.