TONAL PITCH STEP DISTANCE: A SIMILARITY MEASURE FOR CHORD PROGRESSIONS

W. Bas de Haas, Remco C. Veltkamp, Frans Wiering
Departement of Information and Computing Sciences, Utrecht University
{Bas.deHaas, Remco.Veltkamp, Frans.Wiering}@cs.uu.nl

ABSTRACT

The computational analysis of musical harmony has received a lot of attention the last decades. Although it is widely recognized that extracting symbolic chord labels from music yields useful abstractions, and the number of chord labeling algorithms for symbolic and audio data is steadily growing, surprisingly little effort has been put into comparing sequences of chord labels.

This study presents and tests a new distance function that measures the difference between chord progressions. The presented distance function is based on Lerdahl's Tonal Pitch Space [10]. It compares the harmonic changes of two sequences of chord labels over time. This distance, named the Tonal Pitch Step Distance (TPSD), is shown to be effective for retrieving similar jazz standards found in the Real Book [3]. The TPSD matches the human intuitions about harmonic similarity which is demonstrated on a set of blues variations.

1 INTRODUCTION

Among musicians and music researchers, harmony is considered a fundamental aspect of western tonal music. For centuries, analysis of harmony has aided composers and performers in understanding the tonal structure of music. The chord structure of a piece alone reveals modulations, tonal ambiguities, tension and release patterns, and song structure [11]. Not surprisingly, the modeling of tonality and computational harmonic analysis have become important areas of interest in music research. Such models can play an important role in content based music information retrieval (MIR). There are obvious benefits in retrieval methods based on harmonic similarity. Melodies are often accompanied by a similar or identical chord progression, or songs may belong to a class with a specific harmonic structure, e.g. blues or rhythm changes, but also cover songs or variations over standard basses in baroque instrumental music could be identified by their harmony.

In this article we present a method for matching two sequences of symbolic chord labels that is based on Lerdahl's Tonal Pitch Space (TPS). Lerdahl [10] developed a formal music-theoretic model that correlates well with data from psychological experiments and unifies the treatment of pitches, chords and keys within a single model. Our proposed method uses this model to create step functions that represent the change of harmonic distance in TPS over time. Two step functions can be efficiently compared using an algorithm designed by Aloupis et al. [2]. Therefore the proposed measure is named the Tonal Pitch Step Distance (TPSD).

Contribution: We introduce a new distance function in the domain of polyphonic music that measures the difference between chord progressions. It is key invariant, independent of the sequences' lengths, allows for partial matching, can be computed efficiently, is based on a cognitive model of tonality, and matches human intuitions about harmonic similarity. We illustrate the soundness of the distance measure by applying it on a set of blues variations. The efficacy for retrieval purposes is demonstrated in an experiment on 388 Real Book [3] songs.

2 RELATED WORK

Theoretical models of tonality have a long tradition in music theory and music research. The first geometric representations of tonality date back at least two centuries. Some authors have investigated the formal mathematical properties of harmonic structures [14], but of particular interest for the current research are the models grounded in data from psychological experiments (see for reviews, [7] [8]). A notable model is Chew's [5] spiral array. The Spiral Array is founded on music-theoretical principles (the Riemannian *Tonnetz*) and, as the name suggests, places pitches, chords and keys on a spiral. Chords are represented as three-dimensional shapes within the spiral. Despite the fact that distances between pitches are incompatible with empirical findings [9] [10], Chew's model has yielded some useful and theoretically interesting algorithms. Another important model is Lerdahl's TPS [10], which correlates reasonably well with Krumhansl's empirical data [9] and matches music-theoretical intuitions. TPS serves as a basis of the distance function here presented and will be explained in the next section.

A related problem that has gained a lot of attention as well is the problem of automatic chord labeling. Chord labeling is the task of finding the right segmentation and la-

(a) octave (root) level:	0												(0)
(b) fifths level:	0							7					(0)
(c) triadic (chord) level:	0				4			7					(0)
(d) diatonic level:	0		2		4	5		7		9		11	(0)
(e) chromatic level:	0	1	2	3	4	5	6	7	8	9	10	11	(0)

Table 1. The basic space of the tonic chord in the key of C Major (C = 0, C# = 1, . . . , B = 11), from Lerdahl [10].

bels for a musical piece. A chord label consists of a chord root, triadic quality, inversion, and extensions (additional chord notes). Nowadays, several algorithms can correctly segment and label approximately 80 percent of a symbolic dataset (see for reviews [20] [16]). Within the audio domain, hidden Markov Models are frequently used for chord label assignment [17] [18]. It is widely accepted that chord information from symbolic or audio data yields a relevant and musicological valid abstraction that can aid in discerning the structure of a piece and making similarity judgments.

Both research areas previously touched upon are important matters when it comes to MIR. Although the value of chord descriptions is generally recognized, and various models about the cognition of tonality are available, surprisingly little research has focused on how similar chord sequences relate to each other. Attempts include string matching [12] as a measure of similarity, and analyzing a chord sequences on the basis of rewrite rules [15] [19]. Mauch [13] analyzed the frequencies of chord classes, chord progression patterns within the Real Book data [3] that is used in this study as well. Still, we argue that similarity of chord sequences is underexposed within the MIR field, which is also evident from the fact that there is no MIREX track for chord progression similarity.

3 TONAL PITCH SPACE

The TPS is a model of tonality that fits human intuitions and is supported by empirical data from psychology [9] [1]. The TPS model can be used to calculate the distances between all possible chords and to predict corresponding tension and release patterns. Although the TPS can be used for defining relationships between chords in different keys, it is more suitable for calculating distances within local harmonies [4]. Therefore our here presented distance measure only utilizes the parts of TPS needed for calculating the chordal distances within a given key (this is motivated in section 4).

The basis of the TPS model is the *basic space* (see Table 1) which comprises five hierarchical levels consisting of pitch class subsets ordered from stable to unstable. Pitch classes are categories that contain all pitches one or more octaves apart. The first and most stable level (a) is the root level, containing only the root of the analyzed chord. The next level (b) adds the fifth of the chord. The third level (c) is the triadic level containing all pitch classes of the chord. The fourth (d) level is the diatonic level consisting of all pitch classes of the diatonic scale of the current key. The last and least stable level (e) is the chromatic level containing all pitch classes. The shape of the basic space of C major strongly resembles Krumhansl and Kessler's [9] C major-key profile. For every chord change, the levels (a-c) must be adapted properly and for a change of key, level d must be adapted. The basic space is hierarchical: if a pitch class is present at a certain level, it is also present at subsequent levels.

The basic spaces of chords can be used to calculate distances between these chords. First, the basic space is set to match the key of the piece (level d). Then the levels (a-c) can be adapted to match the chords to be compared. The distance between two chords is calculated by applying the Chord distance rule. Some examples of calculation are given in Tables 2 and 3.

The proposed distance measure uses a Chord distance rule that is slightly different from the Chord distance rule defined in TPS [10] and is defined as follows:

> CHORD DISTANCE RULE: $d(x, y) = j + k$, where $d(x, y)$ is the distance between chord x and chord y. j is the minimal number of applications of the Circle-of-fifths rule in one direction needed to shift x into y. k is the number of non-common pitch classes in the levels (a-d) within the basic spaces of x and y together divided by two [2]. A pitch class is non-common if it is present in x or y but not in both chords.
>
> CIRCLE-OF-FIFTHS RULE: move the levels (a-c) four steps to the right or four steps to the left (modulo 7) on level d [3].

When plotted geometrically, the distances exhibit a regular pattern combining the diatonic circle of fifths horizontally and the common tone circle vertically (see Figure 1). If the chordal space is extended, it forms a toroidal structure. Because we are interested in the metrical properties of the chord distance rule it is good to observe that it has a maximum of 13 (see Table 4) [4]. This maximum can be obtained, for instance, by calculating the distance between a C major chord and an E chord containing all notes of the chromatic scale.

[1] The TPS is an elaborate model; due to space limitation we have to refer to [10], chapter 2, pages 47 to 59, for a more detailed explanation and additional examples.

[2] This calculation of k deviates from the calculation described in [10]. Lerdahl defined k as the number of non-common pitch classes in the levels (a-d) within the basic space of y compared to those in the basic space of x. This definition has the undesirable side effect of making the distance function non-symmetrical. Our proposed calculation preserves the symmetry of the distance function and yields equal or similar scores.

[3] If the chord root is non-diatonic j receives the maximum penalty of 3.

[4] Chords with a TPS score of 13 are not musically realistic, but it is useful from a computational point of view to observe that the TPS, and hence the TPSD, has a maximum score.

0							$\underline{7}$				
0		$\underline{2}$					$\underline{\mathbf{7}}$				
0		$\underline{2}$		**4**	$\underline{5}$		$\underline{\mathbf{7}}$				$\underline{11}$
0		2		4	5		7		9		11
0	1	2	3	4	5	6	7	8	9	10	11

Table 2. The distance between the C and G7 in the context of a C major key. The bold numbers in the basic space are specific for the C major chord and the underlined numbers are specific for the G7 chord. The only common pitch class is G (7). All other 9 pitch classes in the levels (a-c) are non-common, therefore $k = \frac{9}{2}$. $j = 1$ because the Circle-of-fifths rule is applied once. The total score is: $1 + 4.5 = 5.5$.

		$\underline{\mathbf{2}}$									
		$\underline{\mathbf{2}}$							$\underline{\mathbf{9}}$		
		$\underline{\mathbf{2}}$			$\underline{5}$	**6**			$\underline{\mathbf{9}}$		
	1	2		4	$\underline{5}$	6	7		9		11
0	1	2	3	4	5	6	7	8	9	10	11

Table 3. The distance of a D and a Dm chord in the context of a D major key. The non-common pitch classes are the minor and major thirds (5 and 6), therefore $k = \frac{3}{2} = 1.5$. Note that the minor third does not belong to the D major scale and adds 2 non-common pitch classes. There is no application of the Circle-of-fifths rule ($j = 0$), hence the total score is $1.5 + 0 = 1.5$.

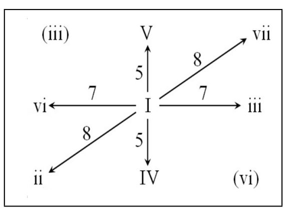

Figure 1. A portion of Lerdahl's chordal space with the roman numbers denoting the seven basic chords in an arbitrary key. The arrows denote the distances as calculated by the Chord distance rule.

0				$\underline{4}$							
0				$\underline{4}$			**7**				$\underline{11}$
$\underline{\mathbf{0}}$	$\underline{1}$	$\underline{2}$	$\underline{3}$	$\underline{\mathbf{4}}$	$\underline{5}$	$\underline{6}$	$\underline{\mathbf{7}}$	$\underline{8}$	$\underline{9}$	$\underline{10}$	$\underline{11}$
0	$\underline{1}$	2	$\underline{3}$	4	5	$\underline{6}$	7	$\underline{8}$	9	$\underline{10}$	11
0	1	2	3	4	5	6	7	8	9	10	11

Table 4. Maximum chordal distance, demonstrated by comparing a C major chord with a chord on E chord containing all pitch classes. The bold numbers are specific for the C major chord and the underlined numbers are specific for the chord on E. The numbers that are underlined and bold are the common pitch classes, the remaining pitch classes are non-common, hence $k = 20$. $j = 3$ because the Circle-of-fifths rule is applied three times and yields its maximal value. The total score is $d(x, y) = \frac{20}{2} + 3 = 13$.

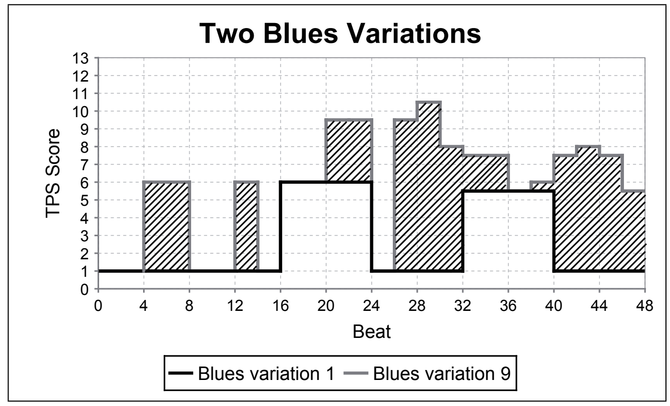

Figure 2. A plot of the step functions of Blues variation 1 and 9 (see Table 5). The TPSD is defined as the hatched area between Blues variation 1 and 9 divided by the length of the shortest step function (which can be either one in this example).

4 TONAL PITCH STEP DISTANCE

On the basis of the Chord distance rule, we define a distance function for chord sequences, called the Tonal Pitch Step Distance (TPSD). The TPSD compares two chord sequences and outputs a number between 0 and the maximum chordal distance 13. A low score indicates two very similar chord sequences and a high score indicates large harmonic differences between two sequences.

The central idea behind the TPSD is to compare the change of chordal distance to the tonic over time. This is done by calculating the chordal distance between each chord of the song and the triadic tonic chord of the key of the song. The reason for doing so, is that if the distance function is based on comparing subsequent chords, the chord distance depends on the exact progression by which that chord was reached. This is undesirable because very similar but not identical chord sequences can then produce radically different scores.

Plotting the chordal distance against the time results in a step function. The difference between two chord sequences can then be defined as the minimal area between the two step functions f and g over all possible horizontal shifts t of f over g (see Figure 2). These shifts are cyclic and to prevent longer sequences from yielding higher scores, the score is normalized by dividing it by the length of the shortest step function. Trivially, the TSPD can handle step functions of different length since the area between non-overlapping parts is always zero.

The calculation of the area between f and g is straightforward. It can be calculated by summing all rectangular strips between f and g, and trivially takes $O(n+m)$ time where n and m are the number of chords in f and g, respectively. An important observation is that if f is shifted along g, a minimum is always obtained when two vertical edges coincide. Consequently, only the shifts of t where two edges coincide have to be considered, yielding $O(nm)$ shifts and a total running time of $O(nm(n+m))$. For the results presented here we used this simple algorithm, but the running time can be further improved to $O(nm \log(n+m))$ by applying an algorithm proposed by Aloupis et al. [2]. They present an algorithm that minimizes the area between two step functions by shifting it horizontally as well as vertically.

We do not use the functionality of TPS to calculate distances between chords in different keys. We choose to do so for two reasons. First, modulation information is rarely present in chord sequence data. Second, we are interested in the harmonic similarity of chord sequences regardless of their keys. This implies that to analyze a chord sequence, the key of the sequence must be provided. The TSPD is not a metric, it does not satisfy the property of *identity of indiscernibles*; since two different chords can have the same Lerdahl distance (see Figure 1), it is possible to construct two different chord sequences with the same TPSD.

5 EXAMPLES

To illustrate how our measure behaves in practice, the distances are calculated for a number of blues progressions. Table 5 shows seventeen variations on a twelve bar blues (the last 12 columns) by Dan Hearle found in [1]. The progressions read from left to right with the numbers in the header denoting the bar. The progression in the top row is a very simple blues and as one moves down the progressions become more complex, more altered, and more difficult to play. The second column displays the distance between the progression in first row and the progression in the current row. The third column shows the distance between two subsequent progressions.

We can make the following observations. The scores correspond well to our intuitive idea of similarity between chord progressions. The scores between similar progressions are small and as progressions become more complex the calculated distance with respect to the most simple blues progression becomes higher.

6 EXPERIMENT

We tested the efficacy of the TPSD for retrieval purposes in an experiment. We used a collection of 388 sequences of chord labels that describe the chords of 242 jazz standards found in the Real Book [3]. The chord label data comes from a collection of user-generated Band-in-a-Box files; Band-in-a-Box is a commercial software package that can be used for generating musical accompaniment. The authors manually checked the files for consistency, quality and correct key. Within this collection, 85 songs contain two or more similar versions, forming 85 classes of songs. These songs have the same title and share a similar melody, but can differ in a number of ways. They can, for instance, differ in key and form, they may differ in the number of repetitions, or have a special introduction or ending. The richness of the chords descriptions can also diverge, i.e. a C7b9b13 may be written instead of a C7, and common substitutions frequently occur. Examples of the latter are relative substitution, i.e. Am instead of C, or tritone substitution, i.e. F#7 instead of C7.

Although additional information about timing and tempo in jazz can contain valuable cues that could be helpful for MIR [6], we only used the chord-per-beat information in our step functions. All songs with multiple versions are used as queries and all other 387 songs are ranked on their TPSD score. We then evaluate the ranking of the other versions.

7 RESULTS

Table 6 shows the results of the experiment. The second and seventh columns display the average first tier per song class. The first tier is the number of correctly retrieved songs within the best $(C-1)$ matches divided by $(C-1)$, where C is the size of the song class. Trivially, by using all songs within a class as a query, C first tiers are calculated and averaged for every class of songs. The third and eighth columns show the average second tier. The second tier is the number of correctly retrieved songs within the best $(2C-1)$ matches, divided by $(C-1)$.

The grand averages [5] of the average first and second tiers are 74 and 77 percent, respectively. This implies that in 74 percent of the song classes the songs searched for are on top of the ranking. Seven song classes in Table 6 contain one or more matches with a TPSD score of 0.00; such a perfect match occurs when two step functions are identical for at least the length of the shortest chord sequence. If these matches are removed from the results, the grand averages of the first and second tiers become 71 and 75 percent, respectively.

8 CONCLUSION

We introduced a new distance function, the Tonal Pitch Step Distance, for chord progressions on the basis of harmonic similarity. The distance of a chord to the tonic triad of its key was determined by using a variant of Lerdahl's Tonal Pitch Space. This cognitive model correlates with empirical data from psychology and matches music-theoretical intuitions. A step function was used to represent the change of chordal distance to its tonic over time. The distance between two chord progressions was defined as the minimal area between two step functions.

[5] Song classes are not weighted on the basis of their size in the calculation of the grand average.

i	d(1,i)	d(i,i-1)	1		2		3		4		5	6		7		8		9		10		11		12	
1	0.00	0.00	F7								Bb7			F7				C7				F7			
2	0.42	0.42	F7								Bb7			F7				C7		Bb7		F7		C7	
3	1.00	0.67	F7		Bb7		F7				Bb7			F7				G7		C7		F7		C7	
4	1.58	0.58	F7		Bb7		F7				Bb7			F7		D7		G7		C7		F7		C7	
5	1.62	0.12	F7		Bb7		F7				Bb7			F7		D7		Gm7		C7		F7		Gm7	C7
6	2.31	0.69	F7		Bb7		F7				Bb7	Eb7		F7		D7		Db7		C7		F7		Db7	C7
7	2.75	1.10	F7		Bb7		F7		Cm7	F7	Bb7	Eb7		F7		Am7	D7	Gm7		C7		Am7	D7	Gm7	C7
8	3.31	0.56	F7		Bb7		F7		Cm7	F7	Bb7	Eb7		Am7		D7		Gm7		C7		Am7	D7	Gm7	C7
9	3.17	0.56	F7		Bb7		F7		Cm7	F7	Bb7	Bm7	E7	F7	E7	Eb7	D7	Gm7		C7	Bb7	Am7	D7	Gm7	C7
10	4.29	2.12	FM7		Em7	A7	Dm7	G7	Cm7	F7	Bb7	Bdim7		Am7	D7	Abm7	Db7	Gm7	C7	Dbm7	Gb7	F7	D7	Gm7	C7
11	5.12	2.08	FM7		Em7	Ebm7	Dm7	Dbm7	Cm7	Cb7	BbM7	Bbm7		Am7		Abm7		Gm7		C7		Am7	Abm7	Gm7	Gb
12	4.88	1.50	FM7		BbM7		Am7	Gm7	Gbm7	Cb7	BbM7	Bbm7		Am7		Abm7		Gm7		Gb7		FM7	Abm7	Gm7	Gb
13	5.23	1.48	FM7		BbM7		Am7	Gm7	Gbm7	Cb7	BbM7	Bbm7	Eb7	AbM7		Abm7	Db7	GbM7		Gm7	C7	Am7	D7	Dbm7	Gb
14	4.40	1.79	FM7		Em7	A7	Dm7	G7	Cm7	F7	BbM7	Bbm7	Eb7	Am7		Abm7	Db7	Gm7		C7		Am7	D7	Gm7	C7
15	4.98	0.75	FM7		Em7	A7	Dm7	G7	Gbm7	Cb7	BbM7	Bm7	E7	Am7		Abm7	Db7	Gm7		C7	Bb7	Am7	D7	Gm7	C7
16	5.42	1.94	F#m7	B7	Em7	A7	Dm7	G7	Cm7	F7	BbM7	Bbm7	Eb7	AbM7		Abm7	Db7	GbM7		Gm7	C7	Am7	D7	Gm7	C7
17	5.71	2.88	FM7		F#m7	B7	EM7	EbM7	DbM7	BM7	BbM7	Bm7	E7	AM7		Am7	D7	GM7		GbM7		FM7	AbM7	GM7	Gb

Table 5. Seventeen blues variations and the TPSD scores between each progression and the first one (second column), and between each progression and the preceding one (third column).

nr.	1st Tier	2nd Tier	Class size	Title of the song class	nr.	1st Tier	2nd Tier	Class size	Title of the song class
1	1.00	1.00	2	A Child is Born	44	1.00	1.00	2	Miyako
2	0.00	0.00	2	A Fine Romance	45	0.50	0.50	2	Moment Notice
3	1.00	1.00	2	A Night In Tunisia	46	1.00	1.00	3	Mood Indigo
4	0.50	0.50	2	All Blues	47	0.73	0.73	6	More I See You, The
5	0.67	0.83	3	All Of Me	48	0.33	0.33	3	My Favorite Things
6	0.50	0.50	2	Angel Eyes	49	1.00	1.00	4	My Funny Valentine
7	0.50	1.00	2	April In Paris	50	1.00	1.00	3	My Romance
8	0.00	0.00	2	Blue in Green	51	0.50	0.50	2	Nefertiti
9	0.50	0.50	4	Corcovado (Quiet Nights of Quiet Stars)	52	0.67	0.83	3	Nica' Dream
10	1.00	1.00	2	Days of Wine and Roses, The	53	1.00	1.00	2	Night Dreamer
11	1.00	1.00	2	Dearly Beloved	54	0.00	0.00	2	Night Has a Thousand Eyes, The
12	0.00	0.00	2	Desafinado	55	0.33	0.33	3	Oleo
13	1.00	1.00	2	Don't Get Around Much Anymore	56	1.00	1.00	3	Once I Loved
14	0.33	0.33	3	Easy to Love	57	1.00	1.00	3	One Note Samba
15	1.00	1.00	2	E.S.P.	58	0.50	0.50	2	Ornithology
16	0.60	0.60	5	Girl from Ipanema	59	1.00	1.00	2	Peace
17	1.00	1.00	2	Green Dolphin Street	60	1.00	1.00	2	Pensativa
18	1.00	1.00	2	Have You Met Miss Jones	61	0.00	0.00	2	Peri's Scope
19	0.70	0.80	5	Here's That Rainy Day	62	1.00	1.00	4	Satin Doll
20	1.00	1.00	4	Hey There	63	1.00	1.00	2	Scrapple from the Apple
21	0.63	0.67	6	How High the Moon	64	0.67	1.00	3	Shadow of Your Smile, The
22	1.00	1.00	3	How Insensitive	65	1.00	1.00	3	Solar
23	1.00	1.00	3	If You Never Come To Me	66	0.33	0.50	3	Some Day My Prince Will Come
24	0.00	0.00	2	I Love You	67	0.33	0.33	3	Song is You, The
25	1.00	1.00	2	I Mean You	68	1.00	1.00	3	Sophisticated Lady
26	1.00	1.00	2	In A Sentimental Mood	69	0.67	1.00	3	So What
27	0.00	0.00	2	Isotope	70	1.00	1.00	3	Stella by Starlight
28	1.00	1.00	3	Jordu	71	1.00	1.00	3	Stompin' at the Savoy
29	1.00	1.00	2	Joy Spring	72	0.67	1.00	3	Straight, No Chaser
30	1.00	1.00	2	Just Friends	73	1.00	1.00	2	Take the "A" Train
31	1.00	1.00	2	Lament	74	1.00	1.00	3	There is No Greater Love
32	1.00	1.00	2	Like Someone In Love	75	0.50	0.50	2	They Can't Take That Away From Me
33	1.00	1.00	3	Limehouse Blues	76	1.00	1.00	2	Triste
34	1.00	1.00	2	Little Waltz	77	1.00	1.00	2	Tune Up
35	0.83	1.00	4	Long Ago and Far Away	78	0.33	0.33	3	Wave
36	1.00	1.00	2	Look to the Sky	79	1.00	1.00	3	We'll Be Together Again
37	0.33	0.33	3	Lucky Southern	80	1.00	1.00	3	Well You Needn't
38	1.00	1.00	3	Lullaby of Birdland	81	1.00	1.00	3	When I Fall In Love
39	1.00	1.00	2	Maiden Voyage	82	0.17	0.25	4	Yesterdays
40	0.33	0.33	3	Meditation	83	0.45	0.60	5	You Are the Sunshine of My Life
41	1.00	1.00	2	Memories of Tommorow	84	1.00	1.00	3	You Are Too Beautiful
42	0.00	0.17	3	Michelle	85	1.00	1.00	2	You Don't Know What Love Is
43	1.00	1.00	2	Misty	**avg.**	**0.74**	**0.77**		

Table 6. 78 song classes of jazz standards found in the Real Book. The first and second tier results show the retrieval efficacy of the TPSD.

We showed that the TPSD as a retrieval method yields promising results. Similar versions of the same jazz standard found in the Real Book can be successfully retrieved on the basis of their chord progressions. The soundness of the TPSD was demonstrated on a series of blues variations.

9 FUTURE WORK

The problem in analyzing the retrieval quality of the TPSD is the lack of good ground truth data. We do not know of any collection of chord sequences that contain user generated similarity labels. It would be interesting to acquire such similarity data and further explore the performance of the TPSD. This could very well be done in a MIREX track. Another issue concerns the used dataset. The Real Book collection is a small collection and only contains songs of a specific musical genre: jazz standards. It would be interesting to explore the performance of the TPSD on a larger dataset and in other musical domains.

We can suggest several improvements of the TPSD as well. Currently, the TPSD does not treat modulations in a musical way. If a modulation occurs within a piece, the scores of the TPSD become very high, and although this enables the TPSD to recognize these modulations, the step function loses nuance. Applying a key finding algorithm locally might yield more subtle results. Another idea is to further exploit the fact that TPSD is very suitable for partial matching by using it for tracing repetitions within a query or matching meaningful segments of a query instead of the query as a whole.

We believe that further improvement of chord labeling algorithms and the development of tools that analyze these labels should be high on the research agenda because chord labels form an abstraction of musical content with substantial explanatory power. In the future we expect TPSD based methods to help users find songs in large databases on the Internet, or in their personal collections. We believe that retrieval on basis of harmonic structure is crucial for the next generation of content based MIR systems.

10 ACKNOWLEDGMENTS

This research was supported by the Dutch ICES/KIS III Bsik project MultimediaN. The authors wish to thank Matthias Mauch (Centre for Digital Music, Queen Mary, University of London) for providing the Real Book data used in this study.

11 REFERENCES

[1] J. Aebersold. *Jazz Handbook*. Jamey Aebersold Jazz, Inc., 2000.

[2] G. Aloupis, T. Fevens, S. Langerman, T. Matsui, A. Mesa, Y. Nuñez, D. Rappaport, and G. Toussaint. Algorithms for Computing Geometric Measures of Melodic Similarity. *Computer Music Journal*, 30(3):67–76, 2004.

[3] Various Authors. *The Real Book*. Hal Leonard Corporation, 6th edition, 2004.

[4] E. Bigand and R. Parncutt. Perceiving Musical Tension in Long Chord Sequences. *Psychological Research*, 62(4):237–254, 1999.

[5] E. Chew. *Towards a Mathematical Model of Tonality*. PhD thesis, Massachusetts Institute of Technology, 2000.

[6] H. Honing and W.B. de Haas. Swing Once More: Relating Timing and Tempo in Expert Jazz Drumming. *Music Perception*, 25(5):471–476, 2008.

[7] C.L. Krumhansl. The Cognition of Tonality - as We Know it Today. *Journal of New Music Research*, 33(3):253–268, 2004.

[8] C.L. Krumhansl. The Geometry of Musical Structure: a Brief Introduction and History. *Computers in Entertainment (CIE)*, 3(4):3–14, 2005.

[9] C.L. Krumhansl and E.J. Kessler. Tracing the Dynamic Changes in Perceived Tonal Organization in a Spatial Representation of Musical Keys. *Psychological Review*, 89(4):334–68, 1982.

[10] F. Lerdahl. *Tonal Pitch Space*. Oxford University Press, 2001.

[11] F. Lerdahl and C.L. Krumhansl. Modeling Tonal Tension. *Music Perception*, 24:329–366, 2007.

[12] C.C. Liu, J. L. Hsu, and A.L.P. Chen. An Approximate String Matching Algorithm for Content-Based Music Data Retrieval. *Proceedings ICMCS*, 1999.

[13] M. Mauch, S. Dixon, C. Harte, M. Casey, and B. Fields. Discovering Chord Idioms Through Beatles And Real Book Songs. *Proceedings ISMIR*, pages 255–258, 2007.

[14] T. Noll and J. Garbers. Harmonic Path Analysis. In Mazzola G. Lluis-Puebla, E. and T Noll, editors, *Perspectives in Mathematical and Computer-Aided Music Theory*. Verlag epOs-Music, Osnabrück, 2004.

[15] F. Pachet. Computer Analysis of Jazz Chord Sequences. Is Solar a Blues? *Readings in Music and Artificial Intelligence*, 1997.

[16] B. Pardo and W.P. Birmingham. Algorithms for Chordal Analysis. *Computer Music Journal*, 26(2):27–49, 2002.

[17] J Pickens. *Harmonic Modeling for Polyphonic Music Retrieval*. PhD thesis, University of Massachusetts Amherst, 2004.

[18] A. Sheh and D.P.W. Ellis. Chord Segmentation and Recognition using EM-Trained Hidden Markov Models. *Proceedings ISMIR*, 3:185–191, 2003.

[19] M.J. Steedman. A Generative Grammar for Jazz Chord Sequences. *Music Perception*, 2(1):52–77, 1984.

[20] D. Temperley. *The Cognition of Basic Musical Structures*. Cambridge, MA, MIT Press, 2001.