POLAR MANHATTAN DISPLACEMENT: MEASURING TONAL DISTANCES BETWEEN CHORDS BASED ON INTERVALLIC CONTENT

Jeff Miller **Johan Pauwels** **Mark Sandler**
Centre for Digital Music
Queen Mary University of London
j.k.miller@qmul.ac.uk j.pauwels@qmul.ac.uk mark.sandler@qmul.ac.uk

ABSTRACT

Large-scale studies of musical harmony are often hampered by lack of suitably labelled data. It would be highly advantageous if an algorithm were able to autonomously describe chords, scales, etc. in a consistent and musically informative way. In this paper, we revisit tonal interval vectors (TIVs), which reveal certain insights as to the interval and tonal nature of pitch class sets. We then describe the qualities and criteria required to comprehensively and consistently measure displacements between TIVs. Next, we present the Polar Manhattan Displacement (PMD), a compound magnitude and phase measure for describing the displacements between pitch class sets in a tonally-informed manner. We end by providing examples of how PMD can be used in automated harmonic sequence analysis over a complex chord vocabulary.

1. INTRODUCTION

Attempts to autonomously label and analyse harmonic sequences in music constitute some of the longest-standing challenges in music information research (MIR) [1]. Various strategies have been applied to chord sequence identification, including information theory [2], graph theory [3–5], and predictive methods such as Hidden Markov Models [6]. Much attention has been given to developing geometric models of musical distance [7–11].

First proposed by Lewin in 1959 [12] and later, in 2007, [13], the discrete Fourier transform can be applied to collections of pitch classes to produce Tonal Interval Vectors (TIVs), which can be used to describe the tonal qualities of chords and pitch class profiles (PCPs) by revealing their constituent intervals and tonal structures [14]. Yust et al. [15] have employed the Fourier transform as a form of cluster analysis on large groups of weighted PCP sets, while Tymoczko and Yust [16] have explored the relationship between voice-leading and Fourier analysis. Other previous work has focused primarily on the magnitude components of TIVs [17, 18], which offer a useful but incomplete picture of intervallic content and tonal quality, as transposition is ignored and certain chord types such as major and minor cannot be disambiguated [19]. Furthermore, many existing methods of measuring distance between TIVs are problematic as they are restricted to pairs of chords and are inconsistent when groups of three or more chords are considered. Additionally, many musical distances falter because they do not capture the directional nature of musical harmonic tension or are adversely affected by enharmonic spellings and conflicting chord vocabularies.

We present the Polar Manhattan Displacement (PMD), a method of describing component-wise directional distance (i.e., displacement) between TIVs which utilises both magnitude and phase information. We demonstrate PMD within the context of the 12-tone equally tempered symbolic domain.

PMD addresses all 4,095 possible pitch class combinations and thus can be applied to any chord vocabulary. PMD offers a consistent displacement measure between chord types (e.g., major7, diminished7, etc.) regardless of transposition or the complexity of the chord vocabulary employed. PMD also measures the intervallic displacement between chords, regardless of chord type. In both cases, displacement measurements are unaffected by transposition of an entire sequence, allowing PMD to identify relative harmonic movements regardless of local key structure.

To demonstrate the utility of PMD, we employ a robust chord vocabulary of 13 chord types including triads, 7th chord types, and suspended chords, as well as chromatic and whole tone scales. We measure displacements amongst these chord types and transpositions, and close by presenting example applications of PMD to automated harmonic analysis and discuss potential applications to other areas of music informatics.

2. BACKGROUND

2.1 Pitch Class Profiles & common musical terms

A pitch class profile (PCP) is a vector of 12 binary values, each representing the categorical presence of its corresponding pitch class in the relevant musical context. This context is set within the time domain; unless otherwise specified, we shall be considering notes which occur si-

Attribution: J. Miller, J. Pauwels, and M. Sandler, “Polar Manhattan Displacement: measuring tonal distances between chords based on intervallic content”, in *Proc. of the 24th Int. Society for Music Information Retrieval Conf.,* Milan, Italy, 2023.

multaneously. Commonly occurring collections of simultaneous pitch events may be referred to as 'chords'. A succession of notes occurring sequentially in ascending or descending pitch order is commonly referred to as a 'scale'. When the time window is increased further and some statistical weighting or filtering is considered, the dominant members of the pitch class set may imply a 'key'.

Throughout this paper, when referring to the collection of possible PCPs, we exclude the empty PCP [0,0,0,0,0,0,0,0,0,0,0,0] which represents an absence of all pitch classes, resulting in $2^{12} - 1 = 4,095$ possible PCPs. To improve clarity, the terms 'chord' and 'scale' shall be considered interchangeable with 'PCP' unless otherwise noted. The term 'chord type' refers to the quality of a chord (e.g., major, minor, etc.) regardless of the chord transposition. A 'chord', however, is a specific combination of root and chord type, such as A_{min} or F_{maj7}.

2.2 Chord vocabulary

For transcription and harmonic analysis purposes, it is useful to focus on the subset of PCPs which correspond to certain chord types. Within the domain of all 4,095 PCPs, the choice of chord vocabulary can be a fairly arbitrary decision. Often, smaller vocabularies of simple chords are chosen to simplify experiments and boost performance scores. There is a risk that an over-simplified chord vocabulary can reduce the usefulness of an analytic system, so it is advantageous that a suitably complex chord vocabulary is employed.

For our examples, we restrict ourselves to a vocabulary of 13 chord types (including, by extension, two scales), but the PMD can be applied to any pitch class profile. Our vocabulary included triads: *major, minor, diminished, augmented, suspended4*; tetrads: *major7, minor7, dominant7, diminished7, half-diminished7, minor/major7*, and scales: *chromatic, wholetone*. Note that some PCPs can be described using different chord types depending on context, thus some chord types, such as maj6, min6 and sus2, are synonymous with other chords already listed in our vocabulary. Regardless of the labels assigned to such synonymous chords, the source PCPs remain the same. For example, PCP [1,0,0,0,1,0,0,1,0,1,0,0] could be described as either C_{maj6} or A_{min7}. Labelling of chords is highly dependent on context and annotator subjectivity. As PMD operates on the basis of underlying PCPs, it is unaffected by such discrepancies in annotation.

2.3 DFTs and Tonal Interval Vectors

By applying a discrete Fourier transform (DFT) to a pitch class profile, we can decompose the PCP into a series of constituent intervallic components. Each of these components will describe the degree to which a particular interval is present within the PCP. Only components $F_1 - F_6$ are required to provide a complete representation, since components $F_7 - F_1 1$ are redundant. Additionally, F_0 reveals the cardinality of the pitch class set; its value is useful for normalisation and allows us to compare any pair of PCPs regardless of the number of pitches present in either.

The resulting 6-dimensional complex vector is referred to as a tonal interval vector (TIV). Scaling may be applied to the various components – often for normalisation purposes, but also in an attempt to more accurately depict perceived cognitive distances between the various interval types within a musical context. [18]

It is worth noting that when generating TIVs, the DFT is applied to the symbolic pitch class vectors and not to audio data. The purpose of the DFT and resulting TIV is to discover the manner in which the pitch classes divide an octave into various musical intervals, and to describe the strength and evenness of each intervallic division.

2.4 Mapping PCPs to TIV space

Each of the 4,095 possible non-empty PCPs produces a unique tonal interval vector. The TIV space therefore is an injection of PCP space: each TIV can be mapped unambiguously to a corresponding pitch class profile. Furthermore, each 6D complex TIV can be represented as a 12-dimensional real-valued vector by converting the complex values of the TIV into magnitude and phase values.

The magnitude and phase values of the recast TIV vector can be represented as a set of 6 tuples (m, p), where m is an unbounded positive real value, and p is a real value $p \in R$ such that $-\pi \leq p \leq \pi$. We will refer to a magnitude and phase tuple (m, p) as a MagPhase tuple. The set of 6 tuples describing Fourier components F_1 - F_6 of a TIV will be referred to as a MagPhase vector.

It has been shown [14,17,18] that converting the Fourier coefficients from complex values to real magnitude and phase values reveals direct correlations to chord type, transposition, and interval ordinality.

2.5 Descriptive properties of Magnitude and Phase

Each TIV Fourier component F_n can be associated with a tonal interval and its complementary inverse interval, e.g. F_1 is associated with the presence of both minor 2nd and major 7th intervals, etc. [14] The magnitude value of a TIV component reflects how strongly the associated interval occurs in the source PCP. For example, F_3 is associated with the presence of major 3rds. Augmented triads (which are composed of nothing but major 3rds) have a maximal F_3 magnitude, whereas diminished 7th chords, which contain no major 3rds, have an F_3 magnitude of 0. Most chords are composed of several interval types and thus have non-zero magnitudes for all 6 components.

Chords containing the same collection of intervals will share a common magnitude profile. This allows some degree of chord classification based on magnitude values. However, magnitude profiles do not convey the ordinal placement of the intervals within the source PCP, meaning that chord types with identical sets of intervals cannot be disambiguated. For example, major and minor triads both contain one of each of the following atomic intervals: $\{m3, M3, P4\}$ and thus will have identical magnitude profiles. For the same reason, magnitude profiles alone do not encode information about the transposition (i.e., the root) of a source PCP, making it impossible to differentiate

C_{dom7} from F_{dom7} or G_{dom7}, for example. By incorporating phase information, both of these shortcomings can be addressed.

3. MOTIVATION

Each component F_1 - F_6 of a TIV describes the strength and position of a particular intervallic quality. By measuring the differences between TIV dimensions separately, the intervallic content of TIVs could be exposed and compared. It is reasonable to consider how distances between them might be useful in describing the relationship between their respective PCPs, as well as enabling computational modelling of musical chords and chord sequences, and the automated study of large corpora.

Distance measures are by definition non-directional. However, harmonic transitions are often asymmetrical in practice, e.g., there is a difference harmonically between a $V7 \rightarrow I$ transition and a $I \rightarrow V7$ transition. By measuring displacement between TIVs, both the distance and direction between them are exposed, which increases the descriptive power of the measure.

To extend the difference measure beyond comparisons of isolated pairs of chords, it would be highly advantageous for the measurements to exhibit collinearity under the addition operation. For example, if two chords of the same type a semitone apart have a particular displacement value, it stands to reason that two chords a whole tone apart should have a displacement double that value. Crucially, this would allow multiple displacements to be summed, making it possible to build sequences and represent various progressions between two chords consistently.

Finally, if differences in chord types and intervallic distances between chords were decoupled, displacement between chords could be visualised on a grid whose orthogonal axes represented each quality. The additive collinearity discussed above would ensure that displacement values could be summed consistently, regardless of which displacements were measured or in what order.

This would allow chord sequences to be represented as traversals across a tonal grid. An example of such a grid is displayed in Figure 1. Chord transitions such as $V7 \rightarrow I$ would have the same displacement regardless of the transposition of the pair. For example, $G_7 \rightarrow C_{maj}$ and $A_7 \rightarrow D_{maj}$ would have an identical displacement, despite being in different keys. Such transpositional invariance would enable an algorithm to catalogue and identify commonly occurring functional sequences.

4. POLAR MANHATTAN DISPLACEMENT

To satisfy the criteria outlined in section 3, we propose the Polar Manhattan Displacement (PMD). PMD measures the component-wise displacement between two TIVs by calculating the differences in the magnitude and phase of each component. As each component F_k describes some aspect of the tonal nature of the source chord, this allows us to measure the displacements between chords in an intervallically informed way.

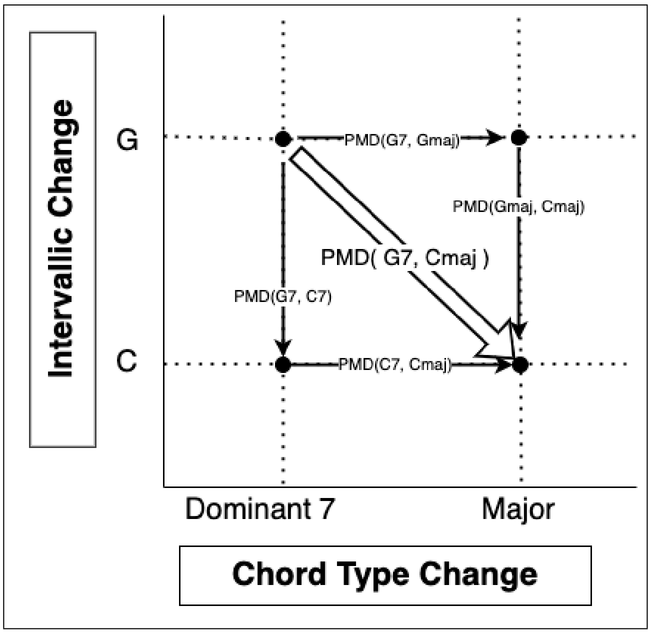

Figure 1. Decoupled chord type and intervallic displacements for the chord transition $G_7 \rightarrow C_{maj}$. There are three routes from $G_7 \rightarrow C_{maj}$; The PMDs of each arrow are summed. All routes result in the same final displacement value. Note that the direction of the progression is significant.

PMD thus borrows from L-1 Manhattan distance but extends it to measure the directional displacements between corresponding magnitudes and phases in each dimension of the polar representation of the TIV plane. The differences are signed, meaning that PMD reflects displacement rather than formal distance. This is advantageous within a musical context because tonal harmony is directional; for example the transition $(C_{maj} \rightarrow G_7)$ has a markedly different tonal impact than $(G_7 \rightarrow C_{maj})$.

4.1 Magnitude processing

The processing of magnitude values is straightforward as it involves only scaling and subtraction. After a TIV is extracted from a PCP, the magnitudes of F_1 to F_6 are divided by the value of F_0. This normalises all magnitudes to the same scale and allows comparison of TIVs having different numbers of pitches in their source PCPs. Further scaling of each component magnitude may be applied to improve the perceptual basis of the space [18] [20]. However, perceptual scaling has not been applied in our study. Following the application of normalisation and scaling, magnitude values are simply subtracted such that the difference between the magnitudes of TIVs U and V is

$$\text{Disp}_{mag}(U \rightarrow V) = (V_{mag} - U_{mag}) \tag{1}$$

4.2 Phase processing

Phase values are simply subtracted in a similar fashion to that presented in section 4.1 such that the angular displace-

Chord type	M1	M2	M3	M4	M5	M6
Maj7						X
Min7		X				X
Dim7	X	X	X		X	X
Aug	X	X		X	X	
Sus4				X		
Chrom	X	X	X	X	X	X
WholeTone	X	X	X	X	X	

Table 1. Some common chord types with zero-magnitude components. 'X' indicates a zero-magnitude vector.

ment between TIVs U and V is

$$\mathrm{Disp}_{phase}(U \rightarrow V) = (V_{phase} - U_{phase}) \tag{2}$$

The cyclic nature of phase information necessitates additional processing. Angular phase values θ need to be cyclically wrapped into the interval $-\pi < \theta \leq \pi$ after all additive and subtractive operations on phase values.

4.3 Definition of PMD

Having defined Disp_{mag} in Eq. (1) and Disp_{phase} in Eq. (2), we can now present the definition of the Polar Manhattan Displacement. PMD is created by concatenating the 6D magnitude and phase displacement vectors Disp_{mag} and Disp_{phase} into one 12D vector.

Given chords (P_j,Q_k) and their corresponding TIVs (U,V),

$$\mathrm{PMD}(P_j, Q_k) = \begin{bmatrix} \mathrm{Disp}_{mag}(U \rightarrow V) \\ \mathrm{Disp}_{phase}(U \rightarrow V) \end{bmatrix} \tag{3}$$

4.4 Zero-magnitude handling

Some PCPs have various TIV components with zero-valued magnitudes. This indicates that the corresponding interval type is absent from the source PCP. While this is of no consequence when calculating magnitude difference (for the magnitude is simply 0), the associated phase of these components is effectively undefined, complicating the calculation of differences between phase values. In our vocabulary, the following chord types contain one or more zero-magnitude component vectors: $\{maj7, min7, dim7, aug, sus4, chrom, wltn\}$. Table 1 details these zero-magnitude chords and their affected components.

We employ the convention that these phase values are considered to be 0. This preserves the additive properties of PMD and allows multiple displacement values to be added together consistently.

4.5 Investigating displacements of type and interval

As indicated in section 3, it would be highly advantageous (and musically interesting) to find a way to distinguish the degree of displacement due to changes in chord types versus shifts of interval. In this section, we examine to what extent such a decoupling is possible.

For convenience, we present the terms Disp_{type} (representing the displacement between chord types) and Disp_{intv} (representing the intervallic displacement between any two chords of the same type). We define each with a functional representation, then provide an example of the relevant function in use. We employ the following convention to represent a movement from one chord to another:

$$(P_j \rightarrow Q_k) \tag{4}$$

where j and k denote chord types from our vocabulary and P and Q denote two arbitrary chord roots. For example, the following represents a movement from G_7 to C_{maj}:

$$(G_7 \rightarrow C_{maj}) \tag{5}$$

where $j = dom7$, $k = maj$, $P = G$, and $Q = C$. Note that the direction of the chord transition is significant.

4.5.1 Type-based displacement

A type-based displacement Disp_{type} can be derived by calculating the PMD of two chords (P_j, Q_k) having the same root (i.e., $P = Q$) but different types ($j \neq k$). Formally,

$$\mathrm{Disp}_{type}(j, k) = \mathrm{PMD}(P_j, P_k) \tag{6}$$

where $j \neq k$.

Significantly, these displacement values are consistent for all pairs of chord types (j, k) regardless of the values of root (P). As an example, the PMD values corresponding to the displacement from a chromatic scale (j) to each of the chords in our vocabulary (k) are detailed in table 2.

4.5.2 Intervallic displacement

Likewise, the intervallic displacement Disp_{intv} between any two chords (P_j, Q_k) should ideally be consistent regardless of the chord types involved. In a similar fashion to the calculation of Disp_{type}, we calculate Disp_{intv} by calculating the PMD of two chords (P_j, Q_k) having different roots (i.e., $P \neq Q$) but identical types ($j = k$).

$$\mathrm{Disp}_{intv}(P, Q, j) = \mathrm{PMD}(P_j, Q_j) \tag{7}$$

where $P \neq Q$. Note that Disp_{intv} is a vector of real number values and is not expressed in semitones.

Crucially, the Disp_{intv} is largely – but not entirely – independent of chord type j. For all chord types j with uniquely non-zero TIV magnitude components, the Disp_{intv} in function of the root interval shift ($P \rightarrow Q$) expressed in semitones is shown in table 3. Chord types j that contain magnitudes of zero (as shown in table 1) have the same PMD as in table 3 for their non-zero components, but both magnitude and phase components of the PMD corresponding to the TIV components with value zero are also zero. By combining tables 1 and 3, the Disp_{intv} for all chord types in our vocabulary can be determined.

Comp	Chrom	Dim7	Wltn	Aug	Min7	Dom7	Maj7	Dim	HDim7	Major	MinMaj7	Minor	Susp4
M_1	0.00	0.00	0.00	0.00	0.18	0.13	0.13	0.33	0.13	0.17	0.25	0.17	0.24
M_2	0.00	0.00	0.00	0.00	0.00	0.25	0.43	0.33	0.25	0.33	0.25	0.33	0.67
M_3	0.00	0.00	0.00	1.00	0.50	0.35	0.71	0.33	0.35	0.75	0.79	0.75	0.33
M_4	0.00	1.00	0.00	0.00	0.50	0.66	0.25	1.00	0.66	0.58	0.25	0.58	0.00
M_5	0.00	0.00	0.00	0.00	0.68	0.48	0.48	0.33	0.48	0.64	0.25	0.64	0.91
M_6	0.00	0.00	1.00	1.00	0.00	0.50	0.00	0.33	0.50	0.33	0.50	0.33	0.33
P_1	0.00	0.00	0.00	0.00	0.52	1.31	0.26	-1.57	-0.26	-2.36	0.00	-1.31	3.14
P_2	0.00	0.00	0.00	0.00	0.00	1.05	0.53	0.00	1.05	0.00	0.00	-1.05	0.00
P_3	0.00	0.00	0.00	0.00	1.57	0.79	0.79	1.57	2.36	0.46	1.25	1.11	0.00
P_4	0.00	0.00	0.00	0.00	-1.05	-1.76	-2.09	0.00	-0.33	-1.57	0.00	-0.52	0.00
P_5	0.00	0.00	0.00	0.00	-0.52	0.26	1.31	-1.57	-1.31	0.79	0.00	-0.26	0.00
P_6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3.14	3.14	3.14

Table 2. PM Displacements for each chord type from our vocabulary as end chord, starting from the chromatic scale and with the same root. Each column is a PMD vector representing the displacement from the chromatic scale chord type. To reverse the direction, invert the signs of the values.

Component	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11	+12
M_1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
M_2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
M_3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
M_4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
M_5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
M_6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P_1	-0.52	-1.05	-1.57	-2.09	-2.62	3.14	2.62	2.09	1.57	1.05	0.52	0.00
P_2	-1.05	-2.09	3.14	2.09	1.05	0.00	-1.05	-2.09	3.14	2.09	1.05	0.00
P_3	-1.57	3.14	1.57	0.00	-1.57	3.14	1.57	0.00	-1.57	3.14	1.57	0.00
P_4	-2.09	2.09	0.00	-2.09	2.09	0.00	-2.09	2.09	0.00	-2.09	2.09	0.00
P_5	-2.62	1.05	-1.57	2.09	-0.52	3.14	0.52	-2.09	1.57	-1.05	2.62	0.00
P_6	3.14	0.00	3.14	0.00	3.14	0.00	3.14	0.00	3.14	0.00	3.14	0.00

Table 3. PM Displacements of each ascending transposition interval in semitones for chord types that have no zero-valued TIV magnitudes. Each column is a PMD vector. Note the symmetry of the column values; intervallic distances are constant, while the sign indicates the direction of transposition. To reverse the direction, (i.e., to transpose down) invert the signs of the values. Also, note that only the phase elements are affected by transposition.

Component	Tritone Substitution		$V_7 - I$	
	(G_7, Db_7)	(A_7, Eb_7)	(G_7, C_{maj})	(B_7, E_{maj})
M_1	0	0	0.043	0.043
M_2	0	0	0.083	0.083
M_3	0	0	0.392	0.392
M_4	0	0	-0.084	-0.084
M_5	0	0	0.161	0.161
M_6	0	0	-0.167	-0.167
P_1	3.142	3.142	0	0
P_2	0	0	0	0
P_3	3.142	3.142	-1.893	-1.893
P_4	0	0	2.285	2.285
P_5	3.142	3.142	0	0
P_6	0	0	3.142	3.142

Table 4. PMD comparisons of a) two tritone substitutions and b) two $V_7 - I$ sequences. Notice that within each pair the PMD values are identical. This indicates that any pair of chords separated by this displacement will constitute a tritone substitution pair or $V_7 - I$ sequence, respectively, regardless of the transposition.

5. PMD EXAMPLES

5.1 Example 1:Tritone substitution

It is well known within musical harmonic practice that certain chords may be substituted for one another to provide alternative or extended versions of an existing or expected harmony. A common example of this is the tritone substitution, wherein a dominant 7th chord can be replaced with a different dominant 7th whose root is a tritone (i.e., an augmented 4th or diminished 5th) away from the original root. The pitches acting as the 3rd and 7th of the original chord are retained, but their functions are swapped. The remaining 2 pitches of the first chord are replaced with other pitches. The overall function is of a new chord that retains the essential character of the original chord, but provides addition harmonic tension.

While issues of perceptual similarity are beyond the scope of the current study, PMD can provide an objective means of numerically describing such relationships. For example, consider the Polar Manhattan Displacements between each of these two tritone substitution pairs: (G_7, Db_7) and $(A_7, E\flat_7)$ as detailed in Table 4. Notice that the PMD values are identical. Any pair of chords separated by this displacement will constitute a tritone substitution pair.

5.2 Example 2: V7 - I detection

There are a number of MIR tasks involving chord estimation, transcription, and automated harmonic analysis that could benefit from the ability to autonomously identify certain chord progressions, particularly those which are harmonically significant. Traditionally, these tasks are hampered by lack of labelled data, inconsistent chord vocabularies, inter-annotator disagreement, etc. As PMD operates on unlabelled symbolic data, it could contribute to addressing such shortcomings and improving performance in automated transcription, labelling, and harmonic analysis. Table 4 describes the displacement between two different $(V_7 \rightarrow I)$ progressions and confirms that PMD can identify and encode the $(V_7 \rightarrow I)$ progression consistently, regardless of key or transpositional context.

6. CONCLUSIONS & FURTHER WORK

We began with a background review of pitch class profiles and some basic terms of musical harmonic structures. We then discussed tonal interval vectors: how discrete Fourier transforms can be applied to PCPs to create TIVs, how TIVs can be represented as vectors containing magnitude and phase values, and how those values describe some aspects of the intervallic construction of a chord. We proposed that it could be useful to measure displacements between these objects, and then described the properties necessary for a robust and self-consistent measure of displacement.

We then presented the Polar Manhattan Displacement, its fundamental components, and the processing required to calculate the measurement of magnitude and phase difference values. There was a brief description of our chord vocabulary and the need for suitable chord vocabularies, and a brief discussion of how to maintain transpositional invariance when dealing with non-existent magnitude vectors.

Having discussed the criteria for a suitable displacement measure, and detailed the functional components of our proposed measurement algorithm, we demonstrated how these components could be aggregated to create the Polar Manhattan Displacement measure. We then provided two examples of potential use cases of PMD, one involving the autonomous identification of tritone substitutions, and the other, $V_7 - I$ progressions.

Future technical work will involve evaluating the robustness of PMD when employed on audio data and at various scales of temporal granularity. It would be interesting to investigate extension of PMD to process non-binary PCPs, such as weighted PCPs and harmonic pitch class profiles. As the additive properties of PMD allow displacements to be summed, we would also like to extend the application of PMD to chord sequence modelling and analysis. Finally, we would like to deploy PMD as part of a large corpus study to investigate chord similarity and harmonic practice.

7. ACKNOWLEDGEMENTS

This research was partially supported by EPSRC grant EP/R512072/1 and the British Broadcasting Corporation through an Industrial CASE studentship in collaboration with the BBC Audio Research Partnership.

8. REFERENCES

[1] T. Fujishima, "Realtime Chord Recognition of Musical Sound: A System Using Common Lisp Music," pp. 464–467, 1999.

[2] M. T. Pearce, "The Construction and Evaluation of Statistical Models of Melodic Structure in Music Perception and Composition," Ph.D. dissertation, City University, London, 2006.

[3] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, "Network Motifs: Simple Building Blocks of Complex Networks," *Science*, vol. 298, pp. 824 – 827, 2002.

[4] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, and U. Alon, "Superfamilies of Evolved and Designed Networks," *Science*, vol. 303, no. March, pp. 1538 – 1542, 2004.

[5] S. Itzkovitz, R. Milo, N. Kashtan, R. Levitt, A. Lahav, and U. R. I. Alon, "Recurring Harmonic Walks and Network Motifs in Western Music," *Advances in Complex Systems*, vol. 9, no. 1 & 2, pp. 121–132, 2006.

[6] C. Harte, M. Sandler, and M. Gasser, "Detecting harmonic change in musical audio," *Proceedings of the*

ACM International Multimedia Conference and Exhibition, pp. 21–26, 2006.

[7] E. Chew, "Towards a mathematical model of tonality," Ph.D. dissertation, 1999. [Online]. Available: http://www-rcf.usc.edu/{~}echew/papers/Dissertation2000/ec-dissertation.pdf

[8] ——, "Out of the Grid and Into the Spiral: Geometric Interpretations of and Comparisons with the Spiral-Array Model," *Computing in Musicology*, vol. 15, pp. 51–72, 2007.

[9] C. Harte, "Towards Automatic Extraction of Harmony Information from Music Signals," Ph.D. dissertation, Queen Mary University London, 2010.

[10] D. Tymoczko, "The geometry of musical chords," *Science*, vol. 313, no. 5783, pp. 72–74, 2006.

[11] ——, *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*. New York, New York, USA: Oxford University Press, 2011.

[12] D. Lewin, "Re : Intervallic Relations between Two Collections of Notes," *Journal of Music Theory*, vol. 3, no. 2, pp. 298–301, 1959.

[13] ——, *Generalized Musical Intervals and Transformations*. Oxford University Press, 2007.

[14] J. D. Harding, "Applications of the Discrete Fourier Transform to Music Analysis," Ph.D. dissertation, Florida State University, 2021.

[15] J. Yust, J. Lee, and E. Pinsky, "A Clustering-Based Approach to Automatic Harmonic Analysis: An Exploratory Study of Harmony and Form in Mozart's Piano Sonatas," *Transactions of the International Society for Music Information Retrieval*, vol. 5, no. 1, pp. 113–128, 2022.

[16] D. Tymoczko and J. Yust, "Fourier Phase and Pitch-Class Sum," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11502 LNAI, pp. 46–58, 2019.

[17] G. Bernardes, D. Cocharro, C. Guedes, and M. E. P. Davies, "Conchord: An Application for Generating Musical Harmony by Navigating in a Perceptually Motivated Tonal Interval Space," *International Symposium on Computer Music Multidisciplinary Research*, pp. 1–16, 2015.

[18] G. Bernardes, D. Cocharro, M. Caetano, C. Guedes, and M. E. Davies, "A multi-level tonal interval space for modelling pitch relatedness and musical consonance," *Journal of New Music Research*, vol. 45, no. 4, pp. 281–294, 2016.

[19] A. Ramires, G. Bernardes, M. E. P. Davies, and X. Serra, "TIV.lib: An open-source library for the tonal description of musical audio," in *Proc. of the 23rd International Conference on Digital Audio Effects (DAFx-20)*, Vienna, 2020, pp. 8–13.

[20] I. Sha'ath, "Estimation of Key in Digital Music Recordings," Master's thesis, Birkbeck College, University of London, 2011.