

Giant Steps In The Interpretation Of A Musical PDP Network

Joshua Hathaway and Michael R. W. Dawson
Department of Psychology, University of Alberta

INTRODUCTION

The key theme of research on musical cognition is discovering the mental representations used to organize and interpret music (Deutsch, 1999; Krumhansl, 1990; Lerdahl & Jackendoff, 1983; Sloboda, 1985). More recently, this approach has been extended to include the study of brain mechanisms related to such representations; that is, the cognitive neuroscience of music (Levitin, 2006; Peretz & Zatorre, 2003). Artificial neural networks are biologically inspired models of cognition (Rumelhart & McClelland, 1986), and interpretations of their internal structure can lead to insights about the nature of representation (Dawson, 2004). Thus, given the rise of both musical cognition and of the cognitive neuroscience of music, it is not surprising that artificial neural networks are frequently used to explore musical cognition (Griffith & Todd, 1999; Todd & Loy, 1991).

An artificial neural network is a computer simulation of a “brain-like” system of interconnected processing units; it generates a desired response to an input stimulus. Activities of a set of input units encode a stimulus. The response of the system is represented as activities in a set of output units. The signal sent by one input processor to an output unit is transmitted through a weighted connection. At the start of a simulation, a network has small, randomly assigned connection weights. The network is then taught by presenting it a set of training patterns, each of which is associated with a known correct response. Each pattern is presented to the network's input units, and on the basis of its existing connection weights, the network generates a response to it. An error term for each output unit is calculated by measuring the difference between the desired response of the unit and its actual response. A learning rule is used to change connection weights in such a way that network error is reduced. Ideally, with repeated presentation of the training patterns, the network learns to generate a correct response to each stimulus.

The current paper describes a musical investigation involving an artificial neural network. We taught a neural network to generate the sequence of chords in a jazz progression called the Coltrane changes. We interpreted the internal structure of this network to reveal an elegant and simple formalism that represents this progression in its entirety. We also used the structure of this network to investigate the complexity of the Coltrane changes in relation to other chord progressions.

We first introduce the notion of chord progressions by describing a particular example (the II-V-I) that is related to the Coltrane changes. Second, we describe

the Coltrane changes using a formalism derived from previous musical investigations with neural networks (Yaremchuk & Dawson, 2005, 2008). Finally, we describe how we trained a neural network to generate the Coltrane changes, how we analyzed its internal structure, and the implications of this interpretation. In particular, we discovered that a network represented transitions between chords in a fashion that could be described in terms of a new musical formalism that we had not envisioned. In short, this paper shows that the interpretation of the internal structure of a musical network can provide new formalisms for representing musical regularities, and can suggest new directions for representational research on musical cognition.

THE II-V-I CHORD PROGRESSION

Most jazz pieces are essentially song structures in which musicians play sequences of chords called chord progressions (Sudnow, 1978). Certain chord progressions are popular because the transition from chord to chord is musically pleasing, and because the progression permits moving from one musical key to another, permitting flexibility (one can start the same progression in any key) and providing musical variety.

This is illustrated in one important chord progression called the II-V-I, a progression that is likely the most commonly encountered in jazz (Levine, 1989). This progression starts in a particular musical key. For example, in the key of C major this progression starts by first playing the D minor seventh chord (Dm7), then by playing the G dominant seventh chord (G7), and ends by playing the C major seventh chord (Cmaj7). In the key of C this is a II-V-I progression because D, the root note of Dm7, is the second note of the C major scale; G, the root note of G7, is the fifth note of the C major scale; and C, the root note of Cmaj7, is the first note of the C major scale.

The II-V-I progression shifts to a new musical key by following the major seventh chord that ends the progression in one key with a minor seventh chord built on the same root that defines the II chord of the next key. For instance, after ending the progression above with Cmaj7, the next chord played is C minor seventh (Cm7), which is the II chord of a new musical key, A#. Table 1 provides a complete II-V-I progression that begins in the key of C major, and ends by returning to this key an octave lower.

When learning to play an instrument, students usually also learn to read musical notation, and are taught the elements of music theory (Martineau, 2008). In

Chord Progression For Key			
Key	II	V	I
C	Dm7	G7	Cmaj7
A#	Cm7	F7	A#maj7
G#	A#m7	D#7	G#maj7
F#	G#m7	C#7	F#maj7
E	F#m7	B7	Emaj7
D	Em7	A7	Dmaj7
C	Dm7	G7	Cmaj7

TABLE I. The II-V-I chord progression, descending from the key of C major. The chords in each row are played in sequence, and after playing one row, the next row is played.

learning such elements, students typically encounter the circle of fifths (Figure 1), if only to determine the number of sharps or flats associated with each musical key signature. Importantly, the circle of fifths captures a great deal of formal structure. For instance, consider a scale written in the key of C (the top of the circle, Figure 1). Scales written in the keys adjacent to C in the circle of fifths (i.e., the keys of F or G) are maximally similar to it, sharing all but one note (Krumhansl, 1990). In other words, the circle of fifths is literally a map in which the more related two keys are, the closer together they are in the circle.

Interestingly, the circle of fifths also maps the root notes of the chords in the II-V-I progression. Consider the chord that starts the Table 1 progression, Dm7. Its root note is D. If one finds that note on the circle of fifths and then moves in a counterclockwise direction, then the next note encountered is the root of the V chord, G7. The next note encountered is C, the root of the I chord, Cmaj7. If one starts again from this final C, the next three notes provide the II-V-I roots for the next key in the progression A#.

Of relevance to the current paper, one can teach the II-V-I progression to an artificial neural network. To do this, a set of input units represents possible musical notes (e.g., each unit represents a particular piano key). The output units represent another set of musical notes in a similar fashion. One presents a chord to the network by turning on the input units that define its component notes. The network is then trained with an error-correcting learning rule (Dawson, 2004, 2005) to turn on the output units that define the notes of the next chord in the progression. Interestingly, an extremely simple artificial neural network called a perceptron (Rosenblatt, 1958) can learn the II-V-I progression of Table 1 (Dawson, 2013). This network has direct connections between input units and output units, and has no hidden units that are the defining property of more powerful artificial neural networks.

The purpose of the current paper is to explore a more complicated chord progression that is an elaboration of the II-V-I progression. This second chord progression is the Coltrane changes, named after its creator, legendary saxophonist John Coltrane (Turner, 1975). We investi-

gate two different questions related to training an artificial neural network to generate the Coltrane changes. First, given that this progression is more complicated than the II-V-I, is a more complex neural network required to generate it? Second, if a neural network can learn to generate the Coltrane changes, then what musical regularities does it represent in its internal structure?

THE COLTRANE CHANGES

John Coltrane introduced the chord progression now known as the Coltrane changes on his seminal 1960 album *Giant Steps*, where it is central to two pieces, “Giant Steps” and “Countdown”. The latter piece paid homage to Miles Davis famous composition “Tune Up” which uses the II-V-I. The Coltrane changes are an elaboration of the II-V-I; it includes the three chords of this older progression, but adds four more chords. Two of these are lead-in chords to the V, and the other two are lead-in chords to the I. Adding additional chords to existing progressions reflects Coltrane’s evolving musical sensitivity. In the liner notes of *Giant Steps* he notes, “I feel like I can’t hear but so much in the ordinary chords we usually have going in the accompaniment.” Table 2 provides the Coltrane changes that reflect the elaboration of the Table 1 progression.

In addition to the particular sequence of chords given in Table 2, we also assumed that chord inversions were used to minimize finger movements (assuming that the progression is being played on a keyboard). Consider the C major 7 tetrachord. In root position, its lowest note is C, followed by the notes E, G, and B. To create the first

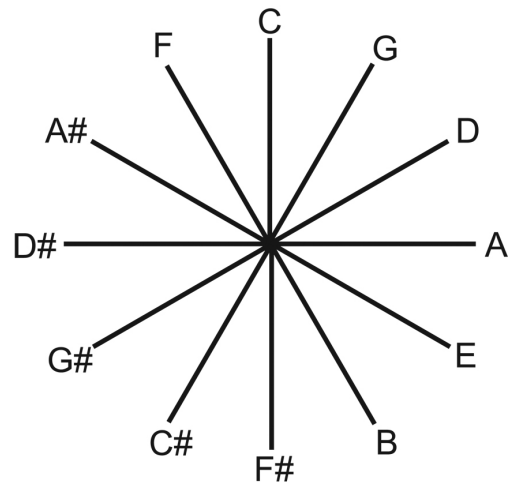


FIG. 1. The circle of fifths. Moving clockwise around the circle, adjacent notes are an interval of a perfect fifth (or seven semitones) apart. Moving counterclockwise around the circle, adjacent notes provide the root notes of chords in the II-V-I progression, as explained in the text.

Chord Progression For Key							
Key	II	Lead-in 1	Lead-in 2	Lead-in 3	Lead-in 4	V	I
C	Dm7	D#7	G#maj7	B7	Emaj7	G7	Cmaj7
A#	Cm7	C#7	F#maj7	A7	Dmaj7	F7	A#maj7
G#	A#m7	B7	Emaj7	G7	Cmaj7	D#7	G#maj7
F#	G#m7	A7	Dmaj7	F7	A#maj7	C#7	F#maj7
E	F#m7	G7	Cmaj7	D#7	G#maj7	B7	Emaj7
D	Em7	F7	A#maj7	C#7	F#maj7	A7	Dmaj7
C	Dm7	D#7	G#maj7	B7	Emaj7	G7	Cmaj7

TABLE II. A progression of Coltrane changes, descending from the key of C major. The chords in each row are played in sequence and after playing one row the next row is played.

inversion of the chord, the C is raised an octave, so that E is now the lowest note. To create the second inversion of the chord, the lowest note of the first inversion (E) is raised an octave, so that G is now the lowest note. To create the third inversion of the chord, the lowest note of the second inversion (G) is raised an octave, so that B is now the lowest note. To minimize finger movements in the Coltrane changes, inversions were applied to the Table 2 chords as follows: the Lead-in 1 chord is played in root position; the II chord, the Lead-in 3 chord and the I chord are played in the first inversion; the Lead-in 2 chord is played in the second inversion; the Lead-in 4 chord and the V chord are played in the third inversion. The use of inversions is evident in the musical score that is presented later as Figure 4.

It is clear that the progression in Table 2 is more complex than the progression in Table 1. Furthermore, the regularities that define the Coltrane changes are notoriously difficult to explain (Capuzzo, 2006; Waters, 2010). In the Giant Steps liner notes, Coltrane describes it as follows: “the bass line is kind of a loping one. It goes from minor thirds to fourths, kind of a lop-sided pattern in contrast to moving strictly in fourths or in half-steps.” Even Coltrane’s accomplished band found his changes hard to master during the rehearsal sessions for the Giant Steps album (Turner, 1975).

To build a training set for teaching the Coltrane changes to an artificial neural network, we discovered a graphical representation of this progression, which is provided in Figure 3 below. This representation takes advantage of the fact that previous interpretations of musical neural networks (Yaremchuk & Dawson, 2008) have revealed musical circles analogous to, but different from, the circle of fifths in Figure 1. Some networks encode musical knowledge using the four circles of major thirds presented in Figure 2. In each of these circles, adjacent notes are a major third (four semitones) apart.

We created a map of the Coltrane changes by placing the circle of fifths in the center, and by then attaching a circle of major thirds to each note of this inner circle, as is shown in Figure 3. (In hindsight, our ability to do this should not have been a surprise, because intervals of major thirds are crucial to the key changes in Coltrane’s progressions (Waters, 2010)). A more elaborate path through this figure (in comparison to the path

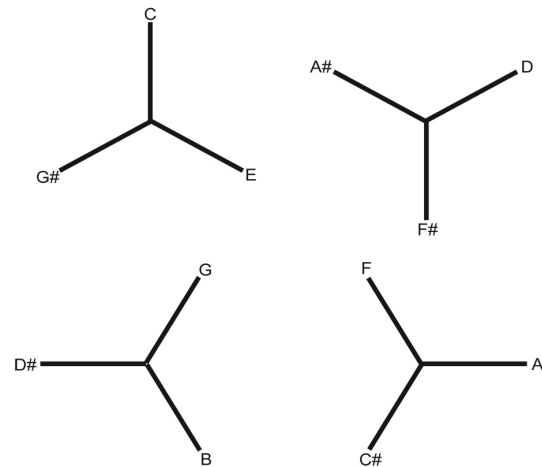


FIG. 2. The four circles of major thirds that Yaremchuk and Dawson (2008) discovered when interpreting the internal structure of an artificial neural network trained to classify musical chords.

through the circle of fifths for the II-V-I) provides the root notes for the seven chords of the Coltrane changes for a given key.

Consider, for example, the key of C. The Coltrane changes begin with Dm7; the root of this chord is in the inner circle of fifths. From this, the next root is the furthest note away on the circle of major thirds attached to the neighboring inner note (i.e. the B attached to the G). The next root is the furthest note away on the circle of major thirds attached to the next note on the inner circle (i.e. the E attached to the C). The progression then returns to the neighboring inner note (the G); the next root is the other outer note attached to it (i.e. the D# attached to the G). Next, the progression returns to the inner C; the next root is the other outer note attached to it (i.e. the G# attached to the C). Finally, the progression moves back to the inner circle of fifths and finished with the root notes G and then C. The Coltrane changes for the next key (A#) begins by keeping on the C, but playing the Cmin7 chord, and then following the identical path from this note to that followed from the D and that was just described. Although this path is complex, it is very easy to grasp when using Figure 3; Figure 3 serves

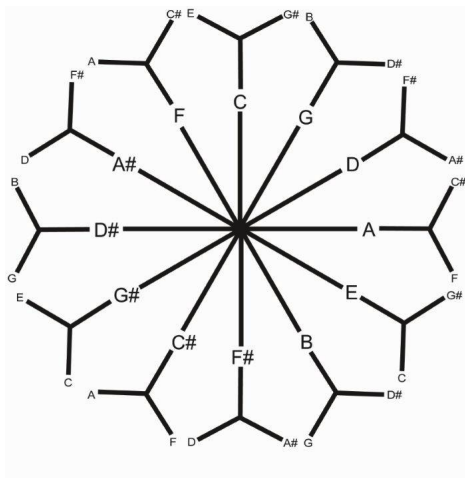


FIG. 3. The Coltrane changes emerge from a map that combines the circle of fifths on the inside with the circles of major thirds on the outside. See text for explanation.

as a musical guide for performing the entire Coltrane changes!

The Coltrane changes are systematic, in the sense that its sequence of root notes lies in a map like Figure 3. However, it seems much more complex than the II-V-I progression described earlier. Clearly, Table 2 involves many more chords than does Table 1, and Figure 3 (and the path through it) is much more complicated than Figure 1 (and the path through it). Is this new progression more complicated from the perspective of music cognition? One way to answer this question is to train a neural network to generate the Coltrane changes. If this progression is indeed more complex from an information processing perspective, then a more complicated network than a perceptron – a network that includes hidden units – is required.

PERCEPTRONS CAN LEARN COLTRANE CHANGES

To begin to answer the complexity question, the first obvious step is to determine whether a simple network, the perceptron, is capable of learning the Coltrane changes. We created a training set to train a perceptron that had 22 input and 22 output unit. Each input and output unit encoded a particular note (i.e. each unit was analogous to a key on a piano). Each output unit used a non-linear, Gaussian activation function, as was the case in previous studies of this type (Yaremchuk & Dawson, 2008). We trained this network in a similar fashion to the method used to train the II-V-I network mentioned earlier: the network was presented a chord in the progression, and an error-correcting rule was used to teach the network to output the next chord in the progression. The entire sequence of chords from Table 2 defined the training set; importantly, the chords were presented in

Inverted Coltrane Changes

M.R.W. Dawson

FIG. 4. The inverted versions of the chords used to train a perceptron to generate the Coltrane changes.

their inverted form (i.e., the notes of the chord were rearranged to produce the most efficient piano fingering); chord inversions are named in the caption of Table 2. Musical notation for the sequence of chords used to train the perceptron is provided in Figure 4.

Given the complexity of the Coltrane changes, we expected that a perceptron would not learn this progression, and that we would then proceed to train a more sophisticated network. To our surprise, the perceptron successfully learned the Coltrane changes in short order. From an information processing perspective, the Coltrane changes are *not* more complex than the II-V-I progression; this is because they both can be learned by artificial neural networks that have no hidden units (i.e. both can be learned by a perceptron).

INTERPRETING THE INTERNAL STRUCTURE OF THE PERCEPTRON

How is it possible for such a simple artificial neural network to represent the complicated sequence of chords that define the Coltrane changes? To answer this question, we carefully examined the weights of the connections from each input unit to each output unit. These weights were highly systematic, and indicated that the network had learned to represent a particular chord as 1) a particular base note (i.e. the lowest note in the chord) and 2) a configuration of finger positions on keys above the base note. By changing the configuration of finger positions (i.e. the positions of the three notes) above the base note, the network could change the type of chord represented along with the chords inversion. Indeed, this makes chord changes possible even when the base note remains the same. For instance, the base note of both D#7 in root position and the base note of G#maj7 in the second version is identical (D# or Eb).

One can move from the first of these chords to the second (as required in the first line in Table 2 above) by keeping the same base note, and by changing the finger configuration above it. The connection weights of the perceptron indicated that the network had literally wired this knowledge in to the connection weights that used input unit activity to turn particular output units on.

The connection weights also revealed that the network had learned that a key aspect of the Coltrane changes was simply encoding the relationships between the different base notes used in the seven chords for the changes in any given key. These relationships only involve four different musical intervals: unison, the major second, the minor 7th, and the major 7th. Each of these intervals was defined by a unique connection weight; that is, the value of a connection weight could be used to read musical interval values from the network.

Furthermore, these relationships only involve three different base notes. Consider beginning the Coltrane changes in the key of C major. The first base note (for the first inversion of Dm7) is F. The next base note is Eb (for the root position of the Lead-in 1 chord); moving from F up to Eb is moving a musical interval of a minor 7th. In the Coltrane changes, the next three chords keep the Eb as the base note; as the note does not change, the musical interval moved is unison. Next, the base note moves on from Eb to F, which is a major second. The progression in this key ends by moving from a base note of F up to a base note of E, which is a major 7th. When the progression begins by playing the base note of the first chord of the next key, the base note shifts from E to Eb, which is another major 7th.

This pattern of shifts from one base note to another, a pattern lifted from the connection weights of the perceptron, can be represented by plotting a series of triangles. The three vertices of one triangle represent the three base notes involved in the seven chords that define the Coltrane changes in a single musical key. Importantly, the triangle that represents these changes in one key, and the triangle that represents the same changes in the next key, share a vertex (i.e., share a base note). Thus one can link the triangles together in a single manifold as is shown in Figure 5.

Figure 5, like Figure 3, provides a musical map through the various base notes of the Coltrane changes. To begin the changes in the key of C, start at F on the lower right of Figure 5. Move next to Eb, which will be repeated as a base note four times. Next, move back to F, which will be used once, and then follow the triangle inside to find that E is the final base note. Moving from E to Eb along the final side of the triangle takes the musician to the base note of the first chord played in the next musical key.

Interestingly, Figure 5 links the Coltrane changes perceptron to the structure of other musical networks studied in our lab (Yaremchuk & Dawson, 2005, 2008). While Figure 5 depicts a ring of intersecting triangles, it also

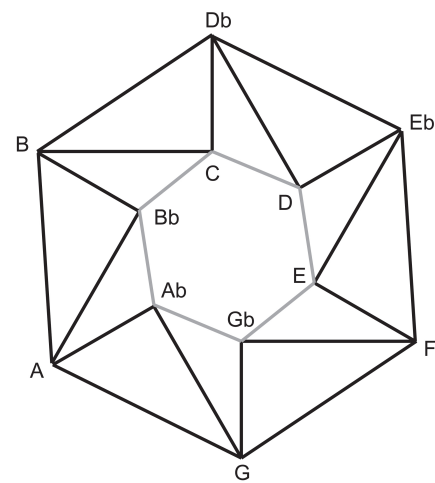


FIG. 5. A map of base note transitions for the Coltrane changes that emerges from examining the connection weights of the trained perceptron. See text for explanation.

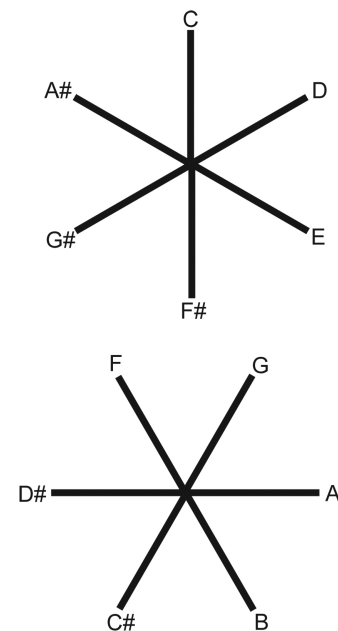


FIG. 6. The two circles of major seconds. The top circle defines the notes around the interior of Figure 5, and the bottom circle defines the notes around the exterior of Figure 5.

can be described as two circles of notes, one on the outside surrounding another on the inside. These two circles are the two circles of major seconds revealed in the connection weights of networks trained to classify chord types; for the sake of completion, the two circles of major seconds are illustrated in Figure 6. In short, the Coltrane changes reflect a specific relation (i.e., a specific orientation in relation to each other) between the two circles of major seconds.

CONCLUSION

The Coltrane changes are taken to be a more complicated elaboration of the classic II-V-I chord progression. The purpose of this paper was to explore the complexity of the Coltrane changes by training an artificial neural network to generate them. To our surprise, we found that an extremely simple network a perceptron learned the Coltrane changes, indicating that they are not more complex than the II-V-I from an information processing perspective. Furthermore, an analysis of the internal structure of this perceptron provided a simple map of base notes used in the progression. This simple map was discovered by the perceptron; the existence of this map permitted such a simple neural network to represent the complexities of the Coltrane changes.

ACKNOWLEDGEMENTS

The research reported in this paper was supported by a SSHRC grant to MRWD.

REFERENCES

- Capuzzo, G. (2006). Pat Martinos The Nature of the Guitar: An intersection of jazz theory and neo-Riemannian theory. *Music Theory Online*, 12(1), 1-17.
- Dawson, M. R. W. (2004). *Minds And Machines: Connectionism And Psychological Modeling*. Malden, MA: Blackwell Pub.
- Dawson, M. R. W. (2005). *Connectionism : A Hands-on Approach* (1st ed.). Oxford, UK ; Malden, MA: Blackwell Pub.
- Dawson, M. R. W. (2013). *Mind, Body, World: Foundations Of Cognitive Science*. Edmonton, AB: Athabasca University Press.
- Griffith, N., & Todd, P. M. (1999). *Musical Networks: Parallel Distributed Perception And Performance*. Cambridge, Mass.: MIT Press.
- Krumhansl, C. L. (1990). *Cognitive Foundations Of Musical Pitch*. New York: Oxford University Press.
- Levine, M. (1989). *The Jazz Piano Book*. Petaluma, CA: Sher Music Co.
- Levitin, D. J. (2006). *This Is Your Brain On Music*. New York, N.Y.: Dutton.
- Martineau, J. (2008). *The Elements Of Music*. New York: Walker & Company.
- Peretz, I., & Zatorre, R. J. (2003). *The Cognitive Neuroscience Of Music*. Oxford ; New York: Oxford University Press.
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386-408.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel Distributed Processing*, V.1. Cambridge, MA: MIT Press.
- Sudnow, D. (1978). *Ways Of The Hand: The Organization Of Improvised Conduct*. Cambridge, Mass.: Harvard University Press.
- Todd, P. M., & Loy, D. G. (1991). *Music And Connectionism*. Cambridge, Mass.: MIT Press.
- Turner, R. (1975). John Coltrane: A biographical sketch. *Black Perspective in Music*, 3(1), 3-16.
- Waters, K. (2010). 'Giant Steps' and the ic4 legacy. *Intégral*, 24, 135-162.
- Yaremchuk, V., & Dawson, M. R. W. (2005). Chord classifications by artificial neural networks revisited: Internal representations of circles of major thirds and minor thirds. *Artificial Neural Networks: Biological Inspirations - Ican 2005, Pt 1, Proceedings*, 3696, 605-610.
- Yaremchuk, V., & Dawson, M. R. W. (2008). Artificial neural networks that classify musical chords. *International Journal of Cognitive Informatics and Natural Intelligence*, 2(3), 22-30.