

# Understanding Musical Discontinuity

Scott Donald Barton

Charlottesville, Virginia

BA, Colgate University, 1998

Master of Music, Brooklyn College Conservatory of Music, 2006

A Dissertation presented to the Graduate Faculty of the University of Virginia

In Candidacy for the Degree of Doctor of Philosophy

University of Virginia

December, 2012

Judith Shatin  
Donald Barton  
Scott Barton  
M. Barton

## Abstract

While our experience of musical (dis)continuity is often powerful and clear, articulating the relations that inspire such percepts is not always easy. Part of the reason for this is that our experience of musical (dis)continuity is influenced by a number of physical, cognitive and perceptual factors, and thus is complex. I will therefore explore (dis)continuity's character by applying ideas from psychology, music theory, algorithmic information theory, and statistics to a variety of musical compositions, including my own. These explorations will describe (dis)continuity as primarily dependent on five key elements: how we bring entities into relations, holistic associations, perceived (dis)similarity of intra- and inter-entity attributes and relational structures (which we can describe in terms of *type* and *degree*); hierarchical organization; and context. These various elements work in isolation and in combination depending on the particular musical situation. Because (dis)continuity is multi-faceted in this way, no singular approach will illuminate the full extent of its richness and complexity. Instead, we must be able to approach (dis)continuity from a number of different perspectives; the one that we choose depends on the specific musical scenario. I will therefore incorporate ideas from the aforementioned disciplines, including structure mapping, transformational distance, and parametric dimensions, to describe, represent and eventually quantify the nature of these characteristics and how they interact to produce (dis)continuity percepts. Such a varied set of tools, which can be used in isolation or in combination, will allow us to describe (dis)continuity in a wide variety of musical styles. This may prove useful not only in identifying commonalities between stylistically diverse musics but also in providing an analytic approach to musics that are resistant to traditional tools. Such flexible, yet rigorous, approach will allow us to illuminate the nature of (dis)continuous relations so that we can analyze and compose (dis)continuous music more thoughtfully.

# Table of Contents

<b>1. Chapter I: Introduction .....</b>	<b>1</b>
<b>1.1. Questions .....</b>	<b>1</b>
<b>1.2. Project.....</b>	<b>2</b>
<b>1.3. Some definitions.....</b>	<b>3</b>
<b>1.4. Some restrictions .....</b>	<b>5</b>
1.4.1. Musical hierarchical level.....	5
1.4.2. Memory and Experience .....	5
<b>1.5. Organization of the Paper .....</b>	<b>6</b>
<b>2. Chapter II: Musical Context.....</b>	<b>8</b>
<b>2.1. Koans – Ted Coffey.....</b>	<b>9</b>
2.1.1. The physical and the perceptual .....	10
2.1.2. Discontinuities in Koans .....	11
2.1.3. Contrast that doesn't create discontinuity in Koans .....	19
2.1.4. Creating interest through discontinuity .....	21
<b>2.2. Summary and Roadmap.....</b>	<b>21</b>
<b>3. Chapter III: Defining musical (dis)continuity.....</b>	<b>24</b>
<b>3.1. Discontinuity .....</b>	<b>24</b>
<b>3.2. Continuity .....</b>	<b>27</b>
3.2.1. (Dis)continuity within whole entities.....	27
3.2.2. (Dis)continuity between distinct entities .....	30
3.2.3. Simultaneous entities.....	42

<b>3.3. Characteristics of continuity relations .....</b>	<b>43</b>
3.3.1. Sameness .....	43
3.3.2. Simplicity .....	44
3.3.3. Regularity, Naturalness.....	49
3.3.4. Regularity / Naturalness and Simplicity.....	51
3.3.5. Continuity via Directionality .....	53
<b>4. Chapter IV: Perceiving musical (dis)similarity .....</b>	<b>62</b>
<b>4.1. Models of Similarity .....</b>	<b>64</b>
4.1.1. Spatial Model .....	65
4.1.2. Featural Model .....	65
4.1.3. Structured Representations .....	67
4.1.4. Synthesis – Structural Transformation .....	77
<b>4.2. (Dis)similarity Perception and context.....</b>	<b>78</b>
4.2.1. Psychological Context .....	79
4.2.2. Musical Context.....	84
<b>5. Chapter V: Parts, Wholes and Hierarchies.....</b>	<b>89</b>
<b>5.1. The identity of parts.....</b>	<b>90</b>
<b>5.2. Perceptual hierarchies.....</b>	<b>92</b>
5.2.1. Overview.....	92
5.2.2. Hierarchies and Attention.....	92
5.2.3. (Dis)continuity on Multiple Hierarchical Levels .....	94
5.2.4. Hierarchies and Context .....	99
5.2.5. Hierarchies and Similarity .....	102
5.2.6. Hierarchies and Transformations.....	103

5.2.7. Varied Saliency Among Hierarchical Levels.....	104
5.2.8. Non-Hierarchies .....	106
<b>6. Chapter VI: Parameters and Dimensions .....</b>	<b>111</b>
6.1. Overview .....	111
6.2. Dimensional representation .....	112
6.2.1. More complex dimensions .....	115
6.2.2. Identifying parametric dimensions.....	119
6.2.3. Parameters that resist dimensional representation.....	119
6.3. Rhythmic parametric change in Planets' O People.....	120
6.4. Interaction of Parameters.....	131
6.4.1. Parameters competing for perceptual saliency .....	131
6.4.2. Parameters combining to form higher-level entities .....	132
6.4.3. Multi-level Parameters .....	132
6.4.4. Parameters creating perceptual structural segmentations.....	133
6.4.5. Research on inter-parametric relationships.....	138
6.5. Associations .....	138
6.6. Approaching Parametrically-Rich Music .....	139
6.7. Summary .....	141
<b>7. Chapter VII: Musical Works .....</b>	<b>144</b>
<b>8. Conclusion and Future Questions.....</b>	<b>188</b>
<b>Appendix A: Other rhythmic dimensions .....</b>	<b>191</b>

# Understanding Musical Discontinuity

Scott Barton

## 1. Chapter I: Introduction

*DISCONTINUITY is a profound musical experience ... the musical experiences that are most memorable are the magical moments when expectation is subverted, when complacency is destroyed, and when a new world opens. (Kramer, 1978, p. 177)*

### 1.1. Questions

Musical discontinuity can be both powerful and puzzling. It can articulate form, effect surprise and (re)direct our attention in striking ways. Because discontinuity is often experientially clear, it is thus surprising when we find it difficult to describe the qualities of such percepts in greater depth.<sup>1</sup> These challenges inspire a number of questions: What is required in terms of content and configuration in order to inspire a sense of (dis)continuity? How does the context that surrounds such juxtapositions affect our percepts? Why do some disjunct pairings work better than others? How could we, as composers, make choices so that we create compelling discontinuities that are governed by generative methodologies other than intuition and randomness? Most of all, *how can we think about and represent entities and relations that are perceived to be discontinuous?*

As we try to answer these questions, we find that describing relationships between discontinuous musical elements is difficult. Describing what such

---

<sup>1</sup> Perhaps this unwieldy nature explains why the amount of formal discourse on the subject has lagged behind the development of extreme discontinuity in compositional practice.

discontinuities mean in the larger context of a piece is equally as challenging. One may choose to define disjunct entities in terms of simple negation (*entity B* is not *entity A*). Such a perspective, while elucidatory in an abstract and general way, tells us little about specific relational characteristics, and therefore does little to illuminate the sides of this multi-faceted object. Instead, we need to be able to distinguish between different kinds of (dis)continuity. We need to be able to identify the components of a (dis)continuity relation so that we can make such distinctions. Finally, we need to understand the perceptual and contextual processes beyond the materials of musical relations that exert important influence on our experience of musical (dis)continuity.

## 1.2. Project

My goal in this thesis is to describe and represent our perceptual experience of musical (dis)continuity by applying ideas from psychology, music theory, algorithmic information theory, and statistics to a variety of musical compositions, including my own. As a result of such applications, I will describe (dis)continuity as primarily dependent on how we bring entities into relations, which is influenced by (dis)similarity judgments (which we can describe in terms of *type* and *degree* of contrast); holistic associations; hierarchical organization; and context. To arrive at these conclusions, I will consider the music that constitutes (dis)continuities in terms of entities, attributes and relations. From there, I will incorporate ideas from the aforementioned disciplines, including structure mapping, transformational distance, and parametric dimensions, to describe, represent and eventually quantify the nature of these characteristics and how they interact to produce (dis)continuity percepts.

These ideas will allow us to describe (dis)continuity in a wide variety of musical styles. This may prove useful not only in identifying commonalities between stylistically diverse musics but also in providing an analytic approach to musics that are resistant to traditional tools (*i.e.*, set theory will likely tell us little about the meaningful

structural characteristics of mashups). To illustrate this flexibility, I will analyze a diverse set of pieces including Birtwistle's *Carmen Arcadiae Mechanicae Perpetuum*, Naked City's *Snagglepuss*, Girl Talk's *Pump It Up*, Ligeti's *String Quartet No. 2*, Zorn's *Road Runner*, Coffey's *Koans*, Planets' *O People* and Messiaen's *Oiseaux Exotiques*, as well as my own works *for steps that grow when climbed*, *Breeding in Pieces*, *Push for Position* and *From Here to There*.

The heterogeneity of this approach is necessitated by the heterogeneity of the objects under scrutiny: postmodern problems require postmodern solutions. As we build a framework that is accommodating to diversity, we discover that (dis)continuity cannot be defined rigidly or absolutely. Instead, different kinds and strengths of (dis)continuity result from a variety of contextual factors that are physical and psychological in nature. The framework that we build must be sensitive to this fact. At the same time, the framework that we build must not be so flexible that it is timid when asked to make a judgment about a given (dis)continuity. The most effective solution will be some mix of flexibility and rigidity, so that we can handle diverse elements yet still be assertive about our conclusions. Such a framework will enable composers and analysts to make and describe a wide range of discontinuous music more meaningfully.

### 1.3. Some definitions

Throughout this discussion, I will refer to entities, attributes, features, dimensions and parameters. By the term *entity*, I mean a whole. This distinction is abstract and flexible, thus an entity can be a note, a melody, a rhythm or a piece. *Attributes, features, and dimensions* are closer in meaning. Recognizing the issues in treating such terms interchangeably, psychologist W. R. Garner addressed the distinction between them in *Aspects of a Stimulus: Features, Dimensions, and Configurations* (1978). According to Garner, we can describe a stimulus in terms of properties, which we can divide into *component properties* and *wholistic properties*. A component property is



one

that helps define a particular stimulus but ... is not synonymous with the stimulus. It often is thought of as an abstracted property of the stimulus in the sense that the attribute exists without any particular manifestation of it. Thus such things as color, size, form, brightness, and linearity are component properties of a stimulus, thus attributes. (102)

Component properties are called *attributes*, which comprise *features* and *dimensions*.

Dimensions are expressed in terms of some positive value, and alternative values are mutually exclusive. For example, a sound is presented at a certain loudness, but that sound cannot project another value of loudness simultaneously. Dimensions can be either continuous or discrete. A *feature* is an attribute of a stimulus that is either projected or is not projected. The concept of zero further differentiates between a dimension and a feature. When a feature is zero (not projected), "there is no information in the stimulus that indicates the potentiality of the feature" (103). Alternatively, zero in the case of dimensions represents either 1) one of the possible values on the dimension (zero loudness) or 2) that the dimension isn't projected. Because dimensions exist as a value, "the full set of dimensions is knowable from any single stimulus" (103-104). The word *parameter* is defined as "a numerical or other measurable factor forming one of a set that defines a system or sets the conditions of its operation" (New Oxford American Dictionary). These ideas are closely related to the idea of dimension as described by Garner, but in this discourse, we can differentiate between the two by using *parameter* to refer to psychological phenomena while *dimension* refers to the representation of that phenomena. For example, we perceive a *tempo* change (parameter) that we can then represent on an appropriate tempo dimension. Note, not all authors incorporate the above distinctions, so in conveying the original language of referenced sources, there may be some confusing of these terms. I will try to clarify these situations as necessary.

## 1.4. Some restrictions

Such a complex topic requires some restrictions. First, I will focus on dramatic examples of musical discontinuity, such as the examples above. It could be argued that discontinuity is present, in some form or another, in almost all musical works. I expect that these univocal selections will reveal intrinsic properties of discontinuity that might have implications for milder manifestations.

### 1.4.1. Musical hierarchical level

In regard to the hierarchical ‘levels’ that perceptual (dis)continuity can occur on, the focus of chapter V, I will focus primarily on the *intermediate* level, where distinct elements such as notes and percussive hits can join to form pitch collections, chords and rhythms, which can then join to form higher-level groupings such as phrases.<sup>2</sup> Exclusive focus on lower hierarchical levels would preclude consideration of important musical structures such as rhythms, temporally separate pitch collections, chord progressions and phrases. Exclusive focus on higher hierarchical levels, which include longer-duration elements such as hypermeters and movements, would present a different set of problems primarily involving abstracted knowledge that resides in long-term memory. Discontinuity seems to be most salient within a range of a few seconds, which aligns with the durational boundaries of both the intermediate level and short-term memory (STM) (this is a proposition that needs to be explored further). Thus, the intermediate level seems an ideal vantage point from which one can describe individual sonic events, structurally important emergent entities and ultimately, discontinuity.

### 1.4.2. Memory and Experience

Because discontinuity separates entities heard in the present from those heard in

---

<sup>2</sup> The intermediate is thus a larger category that comprises a number of levels: the lowest contains individual notes while higher levels contain groupings like chord progressions.

the past, we should consider how the characteristics of our memory systems influence our conceptions of past events and perception of present ones. We must consider a relation as the product of an 'entity in immediate experience' and a 'remembered entity' rather than as two entities in immediate experience, the latter available via tools such as scores. That is not to say that scores are not helpful in analysis, but because they are non-temporal, they are able to deliver us to any point in a work at any time, and for any duration: they give us immediate and pristine images of any part of a work. In contrast, our memories of past sonic events are not unaltered and pristine.<sup>3</sup> Instead, our memories of prior events either persist, dissipate or transform in accordance with a rich mental network of other memories, concepts, associations and phenomenal experience. How memories of musical moments are created, persist and recalled is a significantly complex topic that is outside the scope of this project. Instead of specifying such matters, we will keep in mind that the characteristics of memory may influence some of the percepts that we describe.

## 1.5. Organization of the Paper

To put this discussion in the context of music, chapter II will discuss discontinuity in a number of musical works, with an extended analysis of Coffey's *Koans*. The subsequent chapters will explore the ideas illuminated by these analyses: Chapter III will investigate 1) the configurations that inspire (dis)continuity percepts and 2) (dis)continuity's identity as defined by its associations. As we explore these topics, a common theme will develop: (dis)continuity judgments primarily depend on (dis)similarity judgments. The focus of chapter IV will therefore be on (dis)similarity perception. Chapter V will then consider perceptual parts, wholes and representational hierarchies. We can then describe the various components of hierarchical (dis)similarity

---

<sup>3</sup> In *Concepts of Musical Unity*, Fred Maus insightfully identifies some of the differences between the experience of listening and analytical practice that identifies structural characteristics that are unavailable to perceptual experience.

and (dis)continuity perception in terms of musical parameters; Chapter VI will explore some of the ways we can do this. Chapter VII will show the preceding ideas in practice in my own compositions. Chapter VIII will bring all of these ideas together and posit questions for future inquiry.

## 2. Chapter II: Musical Context

Discontinuities play a prominent role in a diversity of musical works. The qualities and function of such discontinuities vary depending on the specific musical context that surrounds them. To illustrate this diversity, consider the following descriptions:

The aural experience of Naked City's *Snagglepuss* (Figure 2.1) is something like listening to a person with a very short attention span flip through radio stations: noise rock is followed by funk, which is followed by solo piano, which is followed by free jazz, which is followed by cool jazz. The result is a series of short, disproportionately sized blocks that are characterized by sharp edges and diverse content.

[Click to listen](#)

**Figure 2.1. Naked City's *Snagglepuss*.**

Harrison Birtwistle's *Carmen Arcadiae Mechanicae Perpetuum* also contains improbable juxtapositions: a single, quiet, unwavering note suddenly gives way to the *fortissimo*-articulated power of the full ensemble, which catches itself in a rhythmic pattern that progresses like the motion of a bent wheel. This dramatic contrast, stasis versus jagged motion, governs the piece's organization, and is explored in a variety of forms.

[Click to listen](#)

**Figure 2.2. Birtwistle's *Carmen Arcadiae Mechanicae Perpetuum***

Girl Talk's *Pump It Up* is constituted from a variety of sonic sources (and thus is associated with the term *mashup*) so that in some cases, each element of a rhythmic

phrase comes from a different place: a piano sample is followed by a synth stab, then a bass drum sample, followed by an electric keyboard sonority. A listener perceives dissonance between a clear emergent rhythm and the collection of ambiguously related fragments that it comprises.

[Click to listen](#)

**Figure 2.3. Girl Talk's *Pump It Up***

Despite their diversity, these examples share important characteristics: they bring dramatically contrasting elements into relations within small temporal windows, creating clear discontinuities. By examining such examples where our experience of discontinuity is clear, we can learn something about discontinuity's essence, which we can then apply to more ambiguous examples. In that spirit, the balance of this chapter will explore a work, Ted Coffey's *Koans*, where clear discontinuities abound. Such consideration will reveal many of the physical, contextual and perceptual factors that affect our experience of musical discontinuity. The identification of such elements will illuminate directions that the subsequent chapters of this paper will pursue.

## 2.1. Koans – Ted Coffey

The first minute of Ted Coffey's *Koans* (Figure 2.4) projects a wealth of diverse and glaring discontinuities:

[Click to listen](#)

**Figure 2.4. The first minute of Ted Coffey's *Koans*.**

In order to reveal why such judgments are inspired in us, I will look at the discontinuous moments that constitute the piece in detail. This process will involve

identifying the temporal location of the discontinuities, the musical materials that are involved in such relations, and the contextual and psychological factors that promote such percepts. The diversity and contrast of these particular materials will enable me to reveal many of the key concepts that I will discuss throughout this paper.

### 2.1.1. The physical and the perceptual

When analyzing musical discontinuity in *Koans*, one becomes aware of the distinction between the physical and perceptual components of our experience, and that these components occur at different times. Indeed, our perception of discontinuity lags its physical causes, which is a function of the time it takes for us to make sense of a discontinuous relation. The time it takes for us to make sense of a relation is a function of the characteristics of the relation, the context within which it occurs and the speed at which discontinuities are introduced (in *Koans*, discontinuities are introduced quite rapidly making the descriptive process more difficult than one might initially think given the clarity of contrast in the piece). To illustrate the differences between the physical and perceptual in the analysis of *Koans*, I conducted some listening experiments where I placed a marker every time I felt a significant, clear discontinuity had occurred. The following were the delays (in milliseconds) between the beginning of a physical contrast in the musical surface and my recognition of it: 1300, 638, 710, 1600, 1200, 1034, 929, 983, 1700, 1600 (1170 average, which equates to just longer than a second). Thus, we can refer to a number of time locations when describing discontinuity: the *physical* (when the contrasting sound actually begins), the *reported* (when a listener is able to report a discontinuity through some physical means), and the *perceived / cognized* (the duration between the physical and reported during which the listener is interpreting sonic information and determining whether or not she perceives a musical discontinuity).<sup>4</sup> I

---

<sup>4</sup> The question of the interaction between perception and cognition in this context is indeed a complex one, and is beyond the scope of this paper.

make these distinctions in order to avoid the problematic assumption that qualities in a physical musical surface are the same as the perceptual discontinuities that they eventually inspire. The physical and perceptual are, of course, related and thus can be used to explain each other: we need to consider both if we are to understand our experience of musical discontinuity. The following discussion will incorporate both physical and perceptual descriptions of discontinuities.

### 2.1.2. Discontinuities in Koans

The first discontinuity in Koans occurs at 0:10, attributed to a juxtaposition of the jagged and the smooth. The music preceding the discontinuity is characterized by choppiness (sound components are short and segmented) and wide dynamic variation while the music following the change consists of held tones that evenly crescendo over 2.5 seconds. Thus, we can describe the contrast using musical parameters such as *duration* (short vs. long), *density* (many vs. few) and *amplitude* (varied vs. crescendo). These characteristics contribute to the overall character of the musical gestures, which can be described in terms of concepts such as *jagged* and *smooth*. We can then group the latter distinctions within the larger category *topography*. This example shows how descriptions of discontinuity can inhabit a number of hierarchical “levels”, with higher-level descriptions, such as *smooth*, being produced by lower-level descriptions, such as *amplitude* (these relationships will be the focus of chapter V). We also need to consider how we perceive and represent parametric qualities, and how such qualities affect our perception of hierarchical relationships and (dis)continuity (which is the focus of chapter VI).

The next reported discontinuity occurs at 0:16.4 after a silence of 370ms, the longest silence in the piece so far. While the character of the material on either side of the silence differs, there doesn't appear to be significant change along any particular parameter. In fact, there are many similarities between the two sections in terms of



short, choppy components and varying dynamics. It seems then, that the silence that separates the two gestures is the most influential agent in this discontinuity percept. Either the silence itself causes the discontinuity or the silence accentuates differences between the adjacent materials to create a discontinuity percept that may not have occurred if the gestures were directly juxtaposed. I will consider the role of silence in musical discontinuity in subsequent chapters.

The third discontinuity [0:21.4] is similar to the first in that stasis follows frantic choppiness. The antecedent texture (*entity A*) is a dense collection of abbreviated elements that convey varied amplitudes, wide registral range, ambiguous pitch, complexity (it is difficult to identify individual components) and distortion. The subsequent texture (*entity B*) is simpler, projecting consistent pitch; a long, unbroken duration; gradual, smooth dynamic change (crescendo) and a “clean” timbre (free from distortion). We can sum up the parametric contrasts in the following table (Figure 2.5):

Parameter	<i>entity A</i>	<i>entity B</i>
Density	Many, contrasting (complex)	Few (simple)
Duration	Short	Long
Dynamics	many, varied	Consistent (crescendo)
Rhythmic Config	Ambiguous	None
Register	Wide	Narrow
Pitch	Ambiguous	Clear
Timbre	Distorted	Clean

**Figure 2.5. Table of parametric changes contained in the discontinuity at 0:21.4 in *Koans*.**

Thus, there is change in seven categories, and in most cases the changes are significant. These parameters combine to create the sense that a jagged, energetic section shifts to one that is smooth and unchanging, again supporting the notion that musical elements are operating on a number of hierarchical levels to create this contrast. This example

supports the notion that perceived discontinuity is a function of the number and degree of parametric changes.

A different kind of discontinuity occurs at 38:011, reported 777ms after the physical occurrence, where a short melody follows a string of electronic pulses. One of the reasons that the melody is distinct, and therefore creates a discontinuity, is because there haven't been any clear melodic phrases to this point. It follows that if a gesture is unique relative to preceding material, it is potentially discontinuity-promoting. Context therefore is an important part of discontinuity perception, which is a theme that I will explore throughout the paper.

This discontinuity can also be described in terms of parametric change, as before (Figure 2.6):

Parameter	<i>entity A</i>	<i>entity B</i>
Register	Narrow	m3, higher
Pitch	None	Clear
Melody	No	Yes
Timbre	electronic pulse	Synth

**Figure 2.6. Table of parametric changes contained in the discontinuity at 0:38.011 in *Koans*.**

In this case, the significant parametric changes occur along register, pitch, melody and timbre. Thus, discontinuity is created from fewer parametric changes than the previous examples, which inspires us to question how many parameters are required to affect discontinuity (a question that will be addressed in chapter VI). We must entertain the possibility that because of the uniqueness of *entity B*, fewer parametric changes are required to effect the perception of discontinuity here.

The next discontinuity occurs at 0:38.746, reported interval 876ms, where a series of duck sounds follow the aforementioned melodic phrase. Clearly, the two juxtaposed

entities are highly dissimilar, but this juxtaposition creates a conundrum in regard to descriptive methodology. To illuminate why this is the case, we will first proceed as before, by considering parametric change. A summary of the parametric changes is listed in Figure 2.7:

Parameter	Entity A	entity B
Pitch	Clear	Ambiguous
Melody	Yes	Ambiguous <sup>5</sup>
Timbre	Synth	duck sounds

**Figure 2.7. Table of parametric changes contained in the discontinuity at 0:38.746 in *Koans*.**

Here, there are even fewer parametric changes than the previous discontinuity, yet this discontinuity is as salient as any in the work. Further, two of the three categories, pitch and melody, change somewhat moderately (clear → ambiguous). The only parameter that changes significantly is timbre (synth → duck sounds). As in the previous example, the fact that we haven't heard duck sounds to this point contributes to the salience of the contrast (which supports the power of uniqueness). Is it the case that significant change in one parameter could create such a powerful discontinuity? In considering this question, we realize that descriptions that rely solely on musical parameters, as have been defined to this point, do not convey all that is salient about a contrast and that therefore, our descriptive tools need to be expanded.

This moment shows us how associations can play a meaningful role in discontinuity relations and thus, we must expand our descriptions to include them. Here, associational contrast is among the most significant contributors to our sense of

---

<sup>5</sup> Note, the melodic contrast noted in this table (Entity A projects a clear melody while the melodic content of Entity B is ambiguous) conveys perceptual experience in this musical context. When one considers the melodic character of each entity in isolation, we find the melodic similarity between the two is significant. When heard in the context of the piece though, the juxtaposition of the two entities accentuates certain qualities over others (which are discussed in the text following the table), which in this case, occludes the melodic similarity between the two entities.

discontinuity: *entity A* evokes associations with electronic production and electronic music styles while *entity B* evokes associations with the natural world. Describing and quantifying associational contrast is not as straightforward as some of the other musical parameters previously mentioned. There are a number of reasons for this. First, associations are psychological phenomena that reside further from the physical musical surface than do other psychological phenomena such as pitch. Therefore, because associations depend more on a network of internal knowledge, associations will differ from listener to listener, given different listeners have different networks of internal knowledge. A variety of (dis)continuities, and descriptions of them, can be produced as a result. Second, associations don't lend themselves to dimensional representation as easily as other parameters such as tempo, for what kind of dimension would comprise electronically-produced melodies and duck sounds? Dimensional representation and associations will be discussed in detail in Chapter VI. In any case, the contrast of electronic and natural sounds is perhaps the most salient component of this relation, which is strengthened by the fact that to this point, we have not yet heard any sounds that evoke the "natural" world.

Associations also affect the identity of an entity and the way it communicates with us. *Entity A*, the melodic line, is playful in a self-aware way, which is a function of its brevity, simplicity and timbre. The playfulness of the statement is also a function of context. Because the melody appears amongst material that is generally volatile, serious, and intense, its playfulness is augmented. Further, the juxtaposition of such serious and playful gestures appears as absurd to some degree, which also evokes an association with playfulness. *Entity B* inspires similar associations with playfulness, also largely because of context. The juxtaposition of electronic sounds and duck sounds in this context strikes us as absurd, which, as described above, evokes associations with playfulness. The placement of such a gesture close, but not adjacent, to the aforementioned "serious" textures augments its playfulness. Clearly, context is what

influences these interpretations: if we were listening to a field recording, the same duck sounds would not be interpreted as they are in this scenario. Thus, even though there is discontinuity between the melody and the duck sounds, they also share a similar mode of communication characterized by playfulness, which connect them relative to the preceding material. Another level of discontinuity is retroactively created as a function of this grouping, separating the previous material from the melody-duck sounds group (Again, I will discuss hierarchical levels in discontinuity relations in further detail in chapter V).

A sawtooth-type hum follows the duck sounds [0:39.7], which creates a discontinuity for many of the same reasons stated in the preceding paragraph (although the hum does not evoke the same kind of playfulness that the melody and duck sounds do). While there are more parametric changes present here as compared to the preceding discontinuity, the associative contrast appears to be the most powerful of the group: more than anything, we hear natural sounds followed by an electronic hum. Changes in number of elements, duration, rhythmic configuration, etc. seem to accentuate the associative change rather than effect discontinuity as individual agents. This is perhaps because the sequence directs us to listen on the level of higher-order relationships and associations, rather than low-level parameters, and we are therefore most sensitive to higher-level (associative in this case) changes. The low-level parameters influence these higher-level percepts, but they play more of a subordinate role (see Figure 2.8). Hierarchies and attention will be addressed in chapter V.

<b>Parameter</b>	<i>entity A</i>	<i>entity B</i>
Number of elements	4	1
Element Duration	Short	Long
Rhythmic Config	Isochronous	Continuous
Pitch	Ambiguous	Clear
Melody	Ambiguous	None

Timbre	duck sounds	Sawtooth hum
--------	-------------	--------------

**Figure 2.8. Table of parametric changes contained in the discontinuity at 0:39.7 in *Koans*.**

The sawtooth hum is followed by a loud industrial beat that cycles three times, pauses, and cycles twice more [0:41.247]. Again, this discontinuity percept seems to result from the number and degree of parametric changes, which include dynamics, rhythmic configuration and timbre (see Figure 2.9):

<b>Parameter</b>	<i>entity A</i>	<i>entity B</i>
Number of elements	1	Many
Element Duration	Long	Short
Dynamics	Consistent, medium	Changing, loud
Rhythmic Config	Continuous	Metric
Register	Narrow	Wide
Timbre	sawtooth hum	industrial percussion

**Figure 2.9. Table of parametric changes contained in the discontinuity at 0:41.247 in *Koans*.**

Associative change also contributes to this discontinuity. While the sawtooth hum inspires connections with noisy electronic equipment and perhaps early computer music, the industrial beats are strongly associated with specific kinds of electronic pop styles. If the hum represents computer music, the change is a fairly dramatic example of genre switching (particularly of non-dance → dance). If the hum signifies noisy electronic equipment, the juxtaposition is even more dramatic as a “real world” scene is placed immediately next to a musical one. In the actual listening experience, there is probably some combination of these two associations.

We also perceive discontinuity at 0:41.247 because of contrast in higher level attributes. In regard to one of those attributes, topography, the sawtooth hum is almost

as flat as can be: it contains no pitch, dynamic, timbral or rhythmic change. The industrial beats project a dramatically more jagged texture, a function of wide registral jumps and loud, forceful beats linked by sharp percussive elements. Energy states also vary dramatically as the stasis of the sawtooth hum is contrasted with a loud, aggressive and densely packed phrase. As before, these higher level attributes are a function of the combination of lower-level parameters.

In sum, the discontinuity at 0:41.247 results from parametric change (number of elements, element duration, dynamics, rhythmic configuration, register, and timbre), two kinds of associational contrast (early computer music : electronic pop music; sounds of the world : musical gestures) and sudden displacements within higher-level parameters (topography, energy state). Without committing to the relative contribution of each of the elements, we can, at this point, feel comfortable in saying that they all influence our perception, and that their number and quality correspond to the relative strength of the discontinuity.

The industrial beats are interrupted by a sustained, distorted tone at 0:43.819. This juxtaposition is very similar to the discontinuity at 0:41.247 in that a largely stable sonority is contrasted with a jagged rhythmic passage, thus conveying topographic change. There are a few notable differences in this case. Here, the sustained tone is louder and more aggressive than the sawtooth hum, a function of amplitude and distortion. The present sonority also varies more timbrally than the sawtooth hum, as the distortion transforms into feedback and moves in the stereo field over time. Additionally, there is contrast in the imagined mode of production, an association, as the feedback distortion arouses images of human-played electric guitars and saturated amplifiers versus the sawtooth hum, which appears as the product of machines. Another difference between the two discontinuities is the level of contrast in energy states. At 0:43.819, there is indeed some contrast in energy state between *entity A* and *entity B*: the industrial percussion section conveys a certain kind of energy given short /

quickly occurring elements and its regular, beat-based articulations that is not present in the sustained tone. At the same time, both sections assert a loud and aggressive quality. In sum, despite some differences, one doesn't come away with a feeling that there is significant energy contrast between the two segments (which differs from the energy state contrast in 0:41.247). Together, the discontinuity at 0:43.819 contains the same number, but not the same degree, of parametric changes as 0:41.247 (see Figure 2.10). Indeed, this discontinuity seems relatively less potent, which speaks to the importance of the *degree* of changes.

Parameters	<i>entity A</i>	<i>entity B</i>
Number of elements	Many	1
Element Duration	Short	Long
Dynamics	changing, loud	Consistent, loud
Rhythmic Config	Metric	Continuous
Register	Wide	Medium
Timbre	industrial percussion	Distortion / feedback

**Figure 2.10. Table of parametric changes contained in the discontinuity at 0:43.819 in *Koans*.**

### 2.1.3. Contrast that doesn't create discontinuity in *Koans*

While contrast is a major component of (dis)continuity, the two ideas are not the same thing, for even when elements differ significantly in content, character and even source, we can in some cases group them into a single, continuous voice. The first ten seconds of *Koans*, which consists of glitch elements, reverberant synth sounds, short glissandi, speed-manipulated phrases, distortion, snare drum sounds, etc., exemplifies this phenomenon. Even though we realize that many of these individual elements differ in character and originated from different sources (such as the reverberant synth sound and the snare drum sound), they somehow cohere as if they were coming from one voice. Perhaps this percept occurs because there is consistency in the other



communication parameters: elements are of varied lengths, but they are generally short, and organization is generally characterized by quick and contrasting juxtapositions of mercurial and diverse materials. It thus seems that at some rate, given the right kinds of constituent characteristics, significant discontinuities are no longer highlighted, rather, they group together to create an overall sense of jaggedness. As a result, we ignore some dissimilarities and judge these diverse elements as originating from a single, virtual source.

There is also a fairly significant change at 43.314 that *doesn't* inspire a sense of discontinuity. The aforementioned industrial rhythms are interrupted by two bass drum hits. There are many parametric changes at this moment including timbre (the sound of the bass drum in *entity B* is different than the preceding rhythms, which suggests these sounds are coming from a different instrument), volume (*entity B* is louder), and tempo (*entity B* is faster). Despite these many parametric differences, one doesn't perceive a significant discontinuity. One reason for this could be context. If we had heard only the aforementioned industrial rhythms up to this point, then this change would be a significant one and would likely be perceived as discontinuous. Instead, because this juxtaposition occurs in the context of more extreme contrasts that incorporate diverse materials (duck sounds, sawtooth hum, glitch elements, etc.), the similarities between these two moments (both segments are metric, use drum sounds, refer to specific musical genres, have clear pitch, etc.) are accentuated and the moments are subsequently perceived as part of the same gesture. Another reason that we don't perceive discontinuity here could be that two such sonic iterations at the given tempo are not enough to establish a new perceptual grouping in this context. In contrast, a minute-long section comprising such drum hits might inspire us to retroactively perceive a discontinuity at the point of change. Whatever the reasons, it follows that the juxtaposition of two sounds / gestures is not inherently continuous or discontinuous: the perceived (dis)continuity depends on the surrounding context. The context of this

piece is defined by a wide range of possibilities, thus, the most extreme contrasts are necessary to create the perception of discontinuity.

#### **2.1.4. Creating interest through discontinuity**

This piece illustrates how discontinuity can create interest from materials that, by themselves, aren't particularly noteworthy. The electronic pulses → duck sounds → hum / sawtooth waveform sequence exemplifies this phenomenon. The sounds, by themselves, aren't particularly exceptional. When they are juxtaposed though, a context that allows us to perceive higher-order relationships is created. When we listen to the passage, we are listening to the relationships as much (perhaps more) as we are listening to inherent qualities of the sounds themselves. In actuality, we listen to both: we must absorb the intrinsic qualities of the individual sounds and then compare those sonorities to their predecessors, which creates a context and allows the perception of higher-order relationships. This type of compositional organization requires a different kind of listening that offers a new world of potential interest. We can then say that a sawtooth hum sound is interesting not for its intrinsic sonic qualities but for its intrinsic sonic qualities in a particular context. The significant cognitive processing involved in such judgments likely explains why it typically takes a second to report a discontinuity, as mentioned in the beginning of this chapter.

## **2.2. Summary and Roadmap**

The above discussion of *Koans* shows that there are many factors that inspire and contribute to our perception of musical (dis)continuity. The physical characteristics of the music, the perceptual and cognitive interpretation of those characteristics, and the context within which these things appear all affect our musical experience. We quickly see the complexity involved in musical (dis)continuity percepts, which inspires us to

take a step back and consider (dis)continuity as a general concept. The understanding that follows will allow us to explore more specialized topics, and more specific musical examples. This first step, defining (dis)continuity as a musical construct, is the focus of chapter III.

The above analyses support the notion that discontinuity is connected to contrast, which we can describe in terms of (dis)similarity. Indeed, throughout this discourse, a connection between (dis)similarity and (dis)continuity judgments will be made. We must therefore consider how we perceive (dis)similarity between the components of a perceived (dis)continuity, which include entities, attributes and relations. As we try to describe the relative (dis)similarity of different musical entities, we will learn more about how structural characteristics influence our percepts. These topics will be the focus of chapter IV.

*Koans* also shows us the importance of hierarchical distinctions in (dis)continuity judgments. It is usually not the case that we perceive discontinuities because of contrast between two individual elements, rather, we perceive discontinuity when one group of elements contrasts with another (although individual sounds do, in some cases, produce a notion of discontinuity). Thus, if we are to understand why we perceive discontinuity, we need to understand why we group certain elements together, how individual elements relate to each other within a group, and how different groups relate to each other within a hierarchy. By perceiving musical components relative to a hierarchy, we are able to perceive and describe contrast on a variety of levels that differ in terms of distance from the musical surface and degree of abstraction. Hierarchical levels also allow us to explain how we group contrasting elements together into a unified voice. We then can consider which hierarchical levels we favor, and how such tendencies affect our musical experience. These topics are the subject of chapter V.

In regard to describing entities, attributes and relations in more detail, we can use musical parameters, such as pitch, timbre, or rhythmic configuration. The above

analyses highlight a number of issues that we must consider when using parameters to describe perceptual discontinuities. We have to consider what parameters we are going to use to describe a relation. We also have to think about how we are going to define the dimensions that represent those parameters. We then have to consider how we represent musical moments on those parametric dimensions. We also have to recognize that associations, both musical and non-musical, affect our discontinuity percepts, and that the qualities, function and representation of such associations differ from that of traditional musical parameters such as pitch or timbre. These considerations are the focus of chapter VI.

In all cases we realize that context plays a specific role in discontinuity percepts. That is, we cannot do parametric or structural relation analyses in isolation, for whether or not such analyzes are relevant depends on the context in which the musical entities appear. This theme, as previously described, will be discussed throughout the paper.

The preceding provides a roadmap for the explorations of this thesis.

### 3. Chapter III: Defining musical (dis)continuity

#### Argument and Design

As the analyses in the last chapter show, our experience of musical (dis)continuity is a complex topic that is influenced by a number of factors. Some of the latter include the qualities of the entities themselves, their musical and non-musical context, listener knowledge, and our perceptual and cognitive tendencies. When we consider the complexity of each of these categories and then imagine the potential interactions that could occur between them, we start to understand why articulating our experience of musical discontinuity is difficult.

In order to clarify this flexible and complex topic, I will first take a step back in order to view continuity and discontinuity more broadly, after which I will focus more closely on the factors that influence such ideas. The distant perspective will consider how the concepts “discontinuity” and “continuity” are typically defined. From there, I will consider how relational structures affect our sense of (dis)continuity, which will include a discussion of (dis)continuity within whole entities and (dis)continuity between distinct entities. Next, I will reflect on some of the abstract qualities associated with our (dis)continuity percepts, which can help us address the question of what continuity is, what it means, and how it functions beyond its configural characteristics. These associations include sameness, simplicity, regularity, naturalness, and direction. Throughout the chapter, a primary theme will develop: (dis)similarity percepts significantly contribute to our sense of (dis)continuity.

#### 3.1. Discontinuity

Consider the following definitions of the term ‘discontinuity’ from popular reference sources:

1. lack of continuity; irregularity (dictionary.com)
2. Lack of continuity, logical sequence, or cohesion (American Heritage Dictionary / answers.com)
3. The quality or state of being discontinuous; want or failure of continuity or uninterrupted sequence; interrupted condition (oed.com)
4. A break or gap in a structure (oed.com);
5. a break or gap (dictionary.com)

As definitions 4 and 5 convey, musical discontinuity is often produced by breaks and / or gaps, which can be created by tacets, non-silent juxtaposed contrast, surprise and change. For example, we perceive discontinuity when a tacet segments a sonic stream that was previously without such gaps, such as in Figure 3.1:

[Click to listen](#)

**Figure 3.1. Excerpt from Birtwistle's *Carmen Arcadiae Mechanicae Perpetuum*, with gap inserted. Perf. Alarm Will Sound.**

We can't simply equate discontinuity with tacets as many examples of musical patterns that we find to be continuous include gaps, such as Figure 3.2.



**Figure 3.2. Continuous phrase with gaps.<sup>6</sup>**

Discontinuity is also often produced when contrasting entities are juxtaposed, which creates breaks in specific musical parameters. With that said, discontinuity is not simply contrast. The first movement of Bach's *Violin Concerto in A Minor* exhibits breaks via contrast between the solo violin and the larger ensemble, but it is unlikely that most

<sup>6</sup> Examples without any citation were produced by the author.

listeners would describe this music as discontinuous. Often, musical discontinuities result from surprise, or, from a break in an expectational vector. With that said, discontinuity is not simply surprise. The dramatic changes of John Zorn's *Road Runner* continue to evoke notions of discontinuity despite the knowledge and expectation of their arrival (discontinuity persists despite confirmation of an expectational vector). The idea of *change* is obviously important given discontinuity judgments result from some apprehension of sonic change, which can be understood as a sonic break. With that said, not all changes inspire perceived breaks, and not all changes that inspire breaks are perceived as discontinuous. Indeed, our listening intuition supports the position that we don't perceive all sonic changes as discontinuities: change is a quality of music. Discontinuity is not simply change. Thus, it seems that discontinuity is more than just a break or a gap: discontinuity is a break or a gap that occurs in specific contexts.

Because the preceding only partially clarifies discontinuity's meaning, we can, as definitions 1, 2 and 3 do, start at a more basic level by understanding 'discontinuity' as the negation or lack of 'continuity'. Indeed, continuity necessarily precedes discontinuity, that is, a phenomenon must be perceived as continuing before it can be perceived as discontinuing.<sup>7</sup> Conversely, discontinuities demarcate and make clear those phenomena that were previously continuous: we are aware of continuity via distinction, by its edges that are formed by discontinuous neighbors.<sup>8</sup> Importantly, remembered or imagined discontinuity can also make us aware of continuity, such as in the case of a sustained tone that sounds from silence: because we remember the sonic circumstances preceding the tone, we can imagine the tone stopping. As a result, we have a sense of the tone continuing *at the time of occurrence*, rather than just retrospectively after the tone stops. In sum, continuity and discontinuity are defined by

---

<sup>7</sup> This is true even if a perceiver is only consciously aware of a continuity as a result of a discontinuity. The perceiver must have had some non-conscious sense of continuity before the discontinuity occurred: something had to be present for the discontinuity to subsequently interrupt.

<sup>8</sup> One could also argue that discontinuities enable one to make distinctions between simultaneous neighbors. Similarity judgments play an important role in this distinction making process. While the latter is not the focus of this discourse, I will explore these ideas to some extent in chapter II.

each other and, in the case of music, require specific temporal ordering. While this binary formulation helps clarify some aspects of (dis)continuity, we will find that (dis)continuity judgments in real musical experience are not neatly bivalent. Before we tackle that potentially unruly subject, let's take a step back and look more closely at discontinuity's defining other: continuity.

## 3.2. Continuity

*Of material things: The state or quality of being uninterrupted in extent or substance, of having no interstices or breaks; uninterrupted connection of parts; connectedness, unbrokenness*  
(oed.com).

### 3.2.1. (Dis)continuity within whole entities

Continuity describes an entity whose parts are connected without interruption or breaks. Many phenomena in the world demonstrate this characteristic, such as a line:



Or, in music, a sustained pitch (Figure 3.3):

[Click to listen](#)

**Figure 3.3. A Continuous Tone.**

We can perceive such entities as continuous, as having no breaks or interstices, when a notion of sameness, or sufficient similarity, characterizes our perception<sup>9</sup> of them

---

<sup>9</sup> The perceptual part of this equation is important: a judgment of continuity in either of these cases reflects perceptual conclusions rather than physical facts. For instance, the sustained pitch that we hear could potentially be formed by a sequence of discrete pulses occurring at rates above 20 Hz. A being with the ability to distinguish discrete pulses at rates above 20 Hz would not judge such a sequence to be continuous in the same way that a typical human listener would. Similarly, Seurat's *Un dimanche après-midi à l'Île de la Grande Jatte* illuminates the difference between the



at *every* point along a dimension relative to other points (which I will refer to as *uninterrupted continuity*). In the case of a visual line, the dimension is space. That is, if portions of the line in different spatial locations inspire some sense of sameness while at the same time the entity as a whole doesn't contain perceivable interruptions, we perceive the line to be continuous. In the case of sound, the dimension is time. That is, if portions of the sound at different moments in time inspire a sense of sameness, while, at the same time, the listener doesn't perceive or remember interruptions since the sound's inception, we can say the sound is continuous, such as in Figure 3.3.<sup>10</sup>

These conclusions of "sameness" result from judgments of difference. When no difference is detected between two objects, they are the same. When some difference is detected between two objects, they convey some degree of (dis)similarity. Now, one could argue, as Leibniz does via the idea of the Identity of Indiscernibles, that "no two distinct things exactly resemble each other ... no two objects have exactly the same properties" (Forrest, Stanford Encyclopedia of Philosophy). This idea is particularly relevant in a medium that is essentially temporal, such as music. Adopting this perspective, different "parts" of a continuous tone, such as Figure 3.3, are not the same. The part of the tone from 0-100 msec is not the same as the part of the tone from 700-775 msec. The fact that the parts occur at different times means that they are different: they occur in different contexts, they are preceded by different material and they are subject to different receptive responses (such as surprise, or habituation). All of these factors could affect how these different parts are perceived.

---

perceptual and physical in the visual modality. When viewed up close, one sees a scattering of distinct, multi-colored dots. As the viewer distances herself from the painting, the diffuse elements transform into unified and continuous color regions. The lesson in both examples is that continuity is not an inherent property of the physical. Rather, it is a judgment made by perceivers that interact with the physical in specific ways (viewing distance in the case of the painting example). Given that (dis)continuity is essentially a perceptual judgment, I will focus more on perceptual interpretations and less on physical identities. If a listener identifies a sustained note as uninterrupted, we can consider it continuous, regardless of its microscopic phenomenal character. Because we are interested in perceptual interpretations, a key to this discourse will be to explore how we perceptually organize information.

<sup>10</sup> In regard to the lack of interruptions, we must acknowledge that configuration is an equally important part of this equation. That is, if we took the constituent parts of a continuous tone, as in Chapter I, and separated them by a variety of temporal intervals, we would not hear the kind of continuity that is present when we hear the tone in its unbroken form. I will explore the role of configuration in greater depth later in this chapter.

Because we are able to perceive difference within unbroken wholes, then cases of judged identity are better described as similarity comparisons that satisfy the appropriate property-sharing requirements. In the case of a sustained tone, if temporal and contextual characteristics are not among the properties that sameness requires (*i.e.*, sameness doesn't require that different parts of a continuous tone occur in the same context), then we can perceive continuity despite differences in those characteristics. Sameness therefore occurs when *certain*, not *all*, properties are perceived to be common between the parts of an entity. In this paper, the process of comparing two feature sets to illuminate commonality and difference is understood as a similarity comparison. Integrating this idea with Gestalt psychologist Kurt Koffka's description of continuity "a straight line will continue as a straight line ... any curve will proceed in its own natural way, a circle as a circle, an ellipse as an ellipse, and so forth" (153), continuity can be created when terms of required similarity are configured in an un-interrupted way.

When (dis)similarity comparisons do not satisfy the property-sharing requirements of continuity, we perceive discontinuity. We can represent this percept as *entity A*  $\rightarrow$  *entity B*. As *entity B* creates a break or gap in the phenomena, it satisfies the definitions of discontinuity above. A break or gap is typically caused by a significantly contrasting entity (a different sound source; a significant timbral, registral or dynamic shift; silence; a new pitch set, etc.), but as mentioned before, whether or not that contrast produces a notion of discontinuity depends on the type and degree of contrast as well as the musical context. The first two items, type and degree of contrast, will be the focus of chapter VI. The influence of context will be discussed throughout this paper.

It follows that (dis)similarity comparisons significantly influence whether we perceive continuity or discontinuity within whole entities. If similarity is the key to continuity, and thus discontinuity, we must consider what specific musical qualities inspire us to consider two musical entities as similar: this topic will be the focus of chapter IV.

### 3.2.2. (Dis)continuity between distinct entities

Continuity can also exist between entities that we perceive as distinct.

Psychologist Irvin Rock describes this kind of continuity as involving two or more distinct entities or contours and a point or seam where they are joined (33-6). Continuity between distinct entities is common in music, which typically contains a variety of distinct elements. For example, consider Figure 3.4:

[Click to listen](#)

**Figure 3.4. Continuous drum beat consisting of distinct elements.**

In this example, the bass drum, snare drum and tom are temporally and timbrally distinct. Despite such independence, the various entities cohere into a continuous beat. In order to explore how continuity is created in this scenario, we first will accept that we can distinguish auditory entities from their rich sonic surroundings. How we do this, *i.e.*, event distinction on the early processing level, perceptual binding, etc., will not be considered in this paper. The key to understanding this species of continuity then becomes how entities are joined.

We can perceive continuity between distinct entities when we make connections between them, which we do by bringing them into a relation. Importantly, entities, *regardless of their relative similarity*, can constitute a relation. Christopher Hasty, in his article *On the Problem of Succession and Continuity in Twentieth Century Music*, explains:

...the apprehension of difference, far from separating discrete events, requires that we bring events together into a relation, otherwise we could not be aware of difference. This bringing together is the mark of continuity and always involves some sort of overlap since the terms of a relation are mutually dependent and together create a context which could not arise apart from their relation (60).

Thus, events that are distinct from each other, that are in some way unlike, can produce notions of continuity. This should seem controversial, for such a description could also characterize discontinuity, which similarly requires a notion of *difference* that results from bringing events together into a relation. To reconcile this controversy, we can think of discontinuity and continuity as symbiotically related rather than diametrically opposed. Notions of continuity and discontinuity result from the same process of relating, and therefore can be described as poles of a common continuum (hence the term *(dis)continuity*). The distinction between them is a function of how the relation's terms (entities) differ in *type* and *degree* (which I will describe in terms of parametric characteristics in chapter VI). With a sense of difference in terms of *type* and *degree*, we make similarity judgments that govern our sense of connectedness (continuity) or interruption (discontinuity).<sup>11</sup>

Context influences how we characterize both relations and the entities that constitute them. In regard to the latter, our perception of an entity may change when it is brought into a relation: we understand *entity A* differently after *entity B* is put next to it. We must therefore avoid characterizing entities in isolation and then comparing those characterizations to determine relative (dis)similarity and (dis)continuity. Instead, we must consider the characteristics of entities in their relational contexts. Similarly, we must consider relations in their contexts (musical and non-musical), for the latter will also affect whether the difference in *type* and *degree* that characterizes those relations results in judgments of continuity or discontinuity (the correspondence between differences in *type* and *degree* and (dis)continuity percepts is not rigid). We also must keep in mind that context can affect how we perceive a musical hierarchy, the levels of which are formed by the interactions between the various entities and relations of a

---

<sup>11</sup> I am not making a claim of the perceptual causality of this process at this point. In one formulation, we first apprehend fine details of each entity subconsciously, after which we compare each entity on a point-by-point basis, which leads us to a larger sense of (dis)similarity. In the opposite formulation, we are first struck by some sense of (dis)similarity, from which we can then delve further into each entity in order to identify and describe specific low level qualities. Whatever the case may be, my intention here is to point out that there is a relationship between intra-entity qualitative detail, (dis)similarity judgments and notions of (dis)continuity.

musical work (this topic will be explored in chapter V).

The preceding represents the core understanding of (dis)continuity relations that will be applied to musical examples throughout this paper.

### 3.2.2.1. *Temporal Considerations and Intermediaries*

When continuity occurs between distinct entities, a relation is formed whose structure can be described as two temporally-separate entities that are connected by an *intermediary*. The latter, in the context of Rock's formulation, refers to the point or seam where distinct entities are joined. We can describe intermediaries in terms of content and durations. Because the duration of intermediaries specifies the temporal distance between entities, to avoid redundancy, we only need to specify intermediaries when describing such relations.

#### 3.2.2.1.1. *Content of Intermediaries*

The content of a relation's intermediary affects how we perceive (dis)continuity. At the most basic level, intermediaries can be characterized by either *silence* or *sound*.

Intermediaries that contain silence can produce salient discontinuities. Indeed, it seems that we perceptually privilege silence, which could be explained by its prominence as a punctuator of groupings within language and a metaphor for rest between moments of physical action. Whatever the reason, composers have shown an intuitive sense of this power of silence, and have used it widely across styles to demarcate form and create discontinuity. For example, in *Oiseaux Exotiques*, Messiaen demarcates solo piano and orchestral sections with deliberate silences (see Figure 3.5).

[Click to listen](#)

**Figure 3.5. Oliver Messiaen – excerpt from *Oiseaux Exotiques*. Cond. Esa-Pekka Salonen. Perf. Paul Crossley, London Sinfonietta.**

These silences heighten our sense of sectional contrast, which is already significant due to gestural and orchestrational differences. Messiaen thus uses temporal distance to augment qualitative distance, which makes juxtaposed contrast, and subsequently the statement of discontinuity, more clear. Aesthetically, we as listeners appreciate the clarity of this communication and it is one of the reasons that this piece, and this particular form of discontinuous organization, works. In addition, short silences between unlike entities not only make boundaries, and thus formal intent, clear, but also act as sonic palette cleansers. They present us with an amount of time without new information, which allows us to digest the previous sonic statement and gives us an opportunity to prepare for potential statements that could be served next. This use of silence is a common (but not necessary) attribute of successful discontinuous music. Conversely, the lack of palette-cleansing, boundary-making silences could be one reason that we find certain contrasting juxtapositions unsuccessful.

When intermediaries contain *sound*, we must differentiate between the kinds of content that make us perceive a transition between *A* and *B* versus the kinds that inspire us to perceive a separate entity *C* (*A C B*). We can do this by describing the degree of similarity between the component parts of the relation (*entity A*  $\rightarrow$  *intermediary*  $\rightarrow$  *entity B*) relative to specific musical parameters. For example, we can identify and compare the pitch characteristics of *entity A* to that of *entity B*, and then determine how the intermediary's pitch characteristics relate to both entities. Consider Figure 3.6 and Figure 3.7.



[Click to listen](#)

**Figure 3.6. Intermediary contains contrasting pitch content creating three separate entities.**



[Click to listen](#)

**Figure 3.7. Intermediary contains shared pitch content, creating transition between entities A and C.**

The A and C sections of both figures are identical. Pitch-wise, A contains C, C#, D, D# and F, while C contains G, G#, A and A#. The B sections of both figures only differ in terms of pitch: B1's pitch set is a complement to that of A + C (E, F# and B) while B2's pitch set borrows from both A and C. Thus, the pitch content of B2 is considerably more similar to that of A and C than the pitch content of B1. Indeed, when we listen to the two examples, Figure 3.6 sounds more like three distinct entities while Figure 3.7 sounds

more like two distinct entities (*A* and *C*) with a transition between them.

Psychological research that investigates how intermediaries affect continuity perception may also be of use in these descriptions. For example, psychologists Bregman and Dannenbring investigated whether frequency transitions could link a sequence of tones that would otherwise be perceived as separate streams. They found that frequency transitions, which were either full or partial, helped participants determine sequence order versus lack of transitions (Bregman, 134-136). The experiment reinforces the assertion of the preceding paragraph that our perception of transitions and (dis)continuity is a function of the similarity between the content of an intermediary and that of its adjacent entities. This study also speaks to the importance of context: how we make sense of sonic information depends not only on the qualities of *entity A* and *entity B*, but also on their temporal configuration and the material that surrounds those entities. Similar studies could help us describe why certain parametric content inspires us to perceive (dis)continuity.

#### 3.2.2.1.1.1. Interrupted continuity

One important kind of continuity that involves distinct entities and a contrasting intermediary is *interrupted continuity*, which we can represent as (*entity A* → *entity B* → *entity A*). Williams and Colomb use the terms *continuous* and *continual* to highlight this distinction: *continuous* means “without interruption” while *continual* means “an activity continued through time, with interruptions” (23).<sup>12</sup> How is it that *continual* could be equated with *continuous* given the former is defined by interruptions and the latter is defined by lack of interruptions? The answer involves the perception of similarity and memory. Say we perceive two distinct entities that are similar enough to be represented by the same symbol (*‘A’* in this case). We are able to make these judgments because we 1) perceive *‘A2’* now, 2) remember that *‘A1’* occurred at some time in the past and 3) are

---

<sup>12</sup> I will primarily use the term *interrupted continuity*, which is synonymous with *continual*.



able to bring the two percepts into a relation. At the same time, the activity that occurs between the two 'A's must be perceived as sufficiently dissimilar to 'A' to be considered as an other (represented by 'B'), and thus an interruption. These similarity judgments allow a perceiver to connect the two 'A' activities, and the temporally separate periods that contain them, in a way similar to that described in the "continuity within wholes" section. This sense of connectedness is the essence of continuity, so interrupted continuity is indeed a type of continuity despite its violating a fundamental part of continuity's definition.

#### 3.2.2.1.1.1. Continuity Illusion

In the field of auditory perception, interrupted continuity has been studied in the context of the "continuity illusion", which essentially involves a softer sound that is interrupted by a louder sound (often a noise burst), which we can represent as ABA. Listeners, given the right conditions, will judge that the softer sound continues through the louder ones (Bregman, 345; Darwin 2005) (see Figure 3.8).

[Click to listen](#)

**Figure 3.8. Apparent Continuity.** *Bregman, A.S., & Ahad, P. (1996) Demonstrations of auditory scene analysis: The perceptual organization of sound.*

Bregman points out that if discontinuities (which include gaps, attenuations and amplifications) are inserted between A and B (A\_B\_A), then A will not be heard as continuing through B. Bregman attributes this to the fact that we assign a perceptual boundary to A, rather than B, in these cases (350-351).<sup>13</sup> These examples demonstrate how we are able to make perceptual connections between elements that are physically separate. By understanding how we make perceptual connections, we can better

---

<sup>13</sup> When there is no discontinuity between A and B, we place the boundary at the beginning of B, and A therefore continues behind B. If A changes before B, then we attribute the boundary to the end of A, and A is not heard as continuing through B (Bregman, 351).

understand how those connections are broken (as in the A\_B\_A case), creating notions of discontinuity.

The continuity illusion reinforces the idea that the content of intermediaries plays an important role in (dis)continuity perception in at least two ways. First, the insertion of space, or silence into a relation will significantly affect (dis)continuity percepts (in Bregman's words, inserting silence between A and B can move the perceptual boundary from B to A thus creating discontinuity). As discussed in the intermediaries section, in music, the insertion of silence often creates a sense of discontinuity, and this realization has been widely incorporated into compositional practice. Second, the continuity illusion shows how an interrupting sound can preserve the perceived continuity of an interrupted sound given the interrupting sound has the right characteristics. In regard to what these characteristics are exactly, psychologists have primarily manipulated parameters of relatively simple auditory stimuli (such as loudness levels, spectral content and frequency glides). One wonders about what characteristics are necessary to inspire continuity percepts in more complex auditory environments...

Indeed, in the wider world of music, interrupted continuity and the continuity illusion have been explored in a number of contexts. Audio compression algorithms make use of this perceptual characteristic by removing masked information without affecting music's perceived continuity (see Davis Pan's "A Tutorial on MPEG / Audio Compression, 1995, for more details on this process). Audio engineers and producers similarly use dynamics processing in similar ways. In the case of a guitar-sounded chord progression over which a piano melody is introduced, the guitar part can be dynamically attenuated, "ducked", using sidechain compression in concert with the statement of the melody. Done correctly, the melody is heard more clearly and the guitar's perceived dynamic continuity is maintained.

Compositionally, Ives' polytextural *Central Park in the Dark* engages our ability to perceive interrupted continuity over extended durations. The first ~57% of the piece is a

continuous sonic flow voiced by the strings that can be characterized as smooth, slow-moving and dynamically subdued. Over the next minute and a half or so, the orchestra builds to an intense and clamorous peak through the introduction of multiple autonomous musical textures (including interpretations of the sounds of parades, ragtime pianos and street bands) (Ives qtd. in Ballantine, 180). The frenzy comes to an abrupt halt, revealing the smooth, soft strings that were present earlier in the piece (Figure 3.9).

[Click to listen](#)

**Figure 3.9. Charles Ives – excerpt from *Central Park in the Dark*. Cond. Leonard Bernstein. Perf. New York Philharmonic.**

As listeners, we are suddenly struck by the sense that the strings had continued in their understated way during the entire duration of their noisy distracters' interjection. It seems then that we have the capacity to perceive continuity despite relatively long-lasting interrupters. How is this possible? In the case of *Central Park in the Dark*, perhaps the proportions of the sections contribute to our ability to perceive continuity: the "A" section persists for a significant amount of time before the interruption, thus establishing the texture strongly in memory. When the A section returns, our memory of the previous A section is readily evoked, thus creating a palpable connection between the two that subsequently, leads to a sense of continuity. Also necessary in all cases of the continuity illusion is an implicit understanding of masking. That is, we understand that the strings *could have* been sounding during the interjection, but we didn't hear them because they were masked by the louder sounds. *Central Park in the Dark* thus illustrates how continuity can be created in the midst of extreme discontinuity (the latter provided by the polytextural section) over prolonged durations.

### 3.2.2.1.2. Temporal characteristics of Intermediaries

When two entities form a relation, their temporal position, which we can describe in terms of the duration of the intermediary that separates them, also has an important influence on (dis)continuity judgments. A sense of musical continuity requires more than just a sense of coherence or connectedness among its parts. If musical entities are to be perceived as continuous, they must be temporally oriented in specific ways. Along these lines, music theorist C. Catherine Losada describes continuity in relation to *unification*: “The term *unification* implies the existence of cohesive relationships among all parts, the term *continuity* implies a stronger condition where the temporal aspect is essential to the structure of the relationships” (60). What then, are the temporal requirements of (dis)continuity?

A common discontinuity-producing temporal configuration involves short intermediaries. For example, Coffey’s *Koans*, Naked City’s *Snagglepuss* (Figure 3.11) and Zorn’s *Road Runner* (Figure 3.10) all create discontinuity by placing contrasting entities in close temporal proximity with little to no transitional material.

[Click to listen](#)

**Figure 3.10. John Zorn – *Road Runner*. Perf. Guy Klucevsek.**

[Click to listen](#)

**Figure 3.11. Naked City – *Snagglepuss*.**

One of the reasons for this is that when intermediaries are short (or absent), relations are maximally clear. This clarity occurs because: 1) The terms are maximally available as they inhabit immediate experience (*entity B*) and the most recent entity in memory (*entity A*). 2) Terms that are more temporally distant are separated by more interfering and potentially confusing information, diluting the juxtaposition. Because this type of temporal organization allows us to perceive relations clearly, it can produce impactful

discontinuity.

While short intermediaries are often used in discontinuous music, they can also inspire continuity between contrasting entities. Most would consider the sounds in Figure 3.12 as discontinuous:

[Click to listen](#)

**Figure 3.12. The segmented sounds of Girl Talk's *Pump It Up*.**

In the context of Girl Talk's *Pump It Up* though, these disparate sounds are brought together in close proximity to produce a stream of sonic information that perceived as continuous (Figure 3.13):

[Click to listen](#)

**Figure 3.13. Excerpt from Girl Talk's *Pump It Up*.**

There are a few explanations of this percept. First, adjacency is a type of connection in itself, so the elements in Figure 3.13 are more connected than in Figure 3.12, thus they appear as more continuous. Second, and perhaps more importantly, as entities are moved closer together temporally, we are able to combine them into an emergent configuration, in this case, a rhythm, which inspires a sense of continuity between the elements (I will explore this topic in greater depth in chapter V). It follows that when composers juxtapose contrasting entities with short intermediaries, they are, in a sense, balancing continuity and discontinuity, for such organizations could be made even more discontinuous though longer intermediaries. Thus, we cannot say that short intermediaries are an attribute of discontinuous music, for they can produce either continuity or discontinuity percepts, depending on the context.

Longer intermediaries can also promote continuity, or weaken the power of discontinuity-producing contrast, given the right content. For example, consider Mark

Applebaum's *Intellectual Property I*, which "seeks to dramatize a stylistic evolution from one set of aesthetic values to another, in this case from a non-vernacular and minimally determined formal environment to a vernacular one that refers to a specific cultural locus (jazz) and employs a rigid form (12-bar blues)" (Applebaum). The design and affect of the piece is one of gradual and gentle change, which is achieved via the transition's duration and how the content of that transition changes over time.

Transitional durations produce continuity / reduce discontinuity proportionally to their length (the transition here is relatively long). Transitional content produces continuity / reduces discontinuity when it changes in a linear, or easily recognized function, way from *entity A* to *entity B* over time (the path here is generally linear). The reason for this is that linear paths are familiar connectors to us, and thus they help us understand the relationship and connection between contrasting entities, which promotes a sense of continuity. As a result, the music gives us a sense of continuous movement from one type of entity to another, even though if the beginning (*entity A*) and end (*entity B*) were juxtaposed, one would likely be struck by a sense of discontinuity.

One wonders about (dis)continuity's possibilities within other types of distant temporal configurations. Can we perceive two similar entities as continuous if they are separated by a long duration? For example, can we bring an entity heard at the end of a piece into a relation with the memory of an entity that occurred at the beginning of a piece to create a sense of continuity? Many musical forms call for a section or motive to return later in the work. Such repetitions clearly create connections between temporally distant entities, but it is not clear such connections engender a sense of continuity (these moments often feel more like Losada's *unification*). How short does a transition need to be to effect a sense of discontinuity? In what ways do different kinds of content changes over time affect musical discontinuity? These questions would be interesting to explore in both compositional and experimental efforts.

In sum, the temporal characteristics of intermediaries is another factor that affects

the degree to which we consider entities to be (dis)continuous. As the duration of intermediaries and the way that they change over time vary, so too will the types and strengths of (dis)continuity that result. In addition, whether a certain kind of intermediary inspires the perception of discontinuity or not depends on the context surrounding that transition. For example, a 750 msec silent transition may be perceived differently in a piece that contains many 250 msec silent transitions versus a piece that contains many 770 msec silent transitions. Bringing all the influential factors discussed so far together, we can say that (dis)continuity is a function of the relationship between juxtaposed entities (which I will discuss further in chapter IV), the content and temporal characteristics of transitions, higher-level organizations (which I have mentioned so far, but will discuss in more detail in chapter V), and the context within which these various elements function.

### 3.2.3. Simultaneous entities

While the above considers (dis)continuity between distinct, temporally-separate entities, we also must entertain the possibility that discontinuity can be produced by distinct entities that have some overlapping simultaneity. In this scenario, *entity A* is established after which *entity B* is introduced at some later time. As in the previous cases, *entity B* must be sufficiently dissimilar to *entity A* in order to create the notion of discontinuity. Unlike the previous cases though, *entity A* here continues while *entity B* sounds. It is hard to conceive this relation in Rock's terms, for there isn't a seam where one entity ends and another begins. Rather, one entity is superimposed on another to create simultaneous rather than sequential contrast. The beginning of Nancarrow's *Study No. 5* exemplifies this kind of discontinuity. Jerky, low register, staccato chords establish *entity A*. On top of this texture, rapid, bold and contoured runs sound high in the piano's register to establish *entity B*. The two entities are so unlike that one is struck by a sense of discontinuity when *entity B* enters, even though *entity A* continues to sound

during these interruptions. The fact that *entity A* continues seems to be incongruous with the descriptions of discontinuity offered to this point.

To understand simultaneously-occurring entities as discontinuous, we must reformulate our understanding of “interruption”. *Entity B* feels as if it interrupts *entity A* in some way, but identifying it as an interruption is problematic because *entity A* doesn’t actually stop. It seems then, that *entity B* is interrupting something (hence the perception of discontinuity) that is related to *entity A*, but is, in fact, not *entity A*. One potential explanation of this phenomenon is that we have conception of a piece as a whole that is informed by, but separate from, constituent entities. In the beginning of *Study No. 5*, our piece-level conception is filled completely by *entity A*. When contrasting *entity B* is introduced, it interrupts this conception as the character of the piece as a whole differs significantly from its character pre-*entity B*. This disjunction creates a sense of discontinuity, even though *entity A* continues as it had before. I will discuss these kinds of hierarchical (dis)continuities in greater depth in chapter V.

### 3.3. Characteristics of continuity relations

We can also describe the characteristics of continuity relations in holistic ways. These descriptions involve concepts that we associate with continuity that reflect some part of its essence, and thus offer us another way to understand what we mean when we say a collection of elements is continuous. The following will explore some of continuity’s characteristic associations, which include sameness, simplicity, regularity, naturalness and direction. In all cases, these associations will point to the influence of (dis)similarity percepts, which I will explore in chapter IV.

#### 3.3.1. Sameness

As described in the section “(dis)continuity within whole entities”, continuity can be produced when sameness, or sufficient similarity, and un-interruption



characterize an entity. Further, as was described in the case of interrupted continuity, the sense of sameness, or sufficient similarity, between entities can create notions of continuity, even when contrasting material separates those entities. In both cases, our ability to perceive sameness, or sufficient similarity, within or between entities creates a strong feeling of connectedness, which is one of the primary attributes of continuity.

### 3.3.2. Simplicity

When entities and relations proceed in a way that we characterize as simple, we can perceive them as continuous. Gestalt psychologist Max Wertheimer, in explaining the law of the continuing curve, states: “If several possibilities are present by which a part may be continued in a whole, the simpler, more regular will be chosen” (Helson, 28). The connection between simplicity and continuity seems plausible, but in order to articulate it, we need to understand what the words “simplicity” and thus “complexity” actually mean. While the preceding question is exhaustingly broad, well-trodden and highly context-dependent, we can consider a few ideas on the subject that can let us start to scratch the surface of how simplicity and continuity are related in musical contexts.

Consider the following definitions of the concept “simple”. The New Oxford American Dictionary describes the word as “easily understood or done; presenting no difficulty”. Jeff Pressing (1997) approaches complexity (and thus simplicity) via the idea of *generative complexity*, which is the length of the shortest program that can generate the entity in question: the longer the length, the greater the complexity. This idea is related to *information-based complexity*, which is defined as the minimal cost of deriving a solution after identifying relevant variables and describing available operations.

Applying these ideas to music, one could argue that the musical progression that is easiest to understand, would require the shortest generative program and would incur the least cost in deriving a solution, is a repeat, as in Figure 3.14.



**Figure 3.14. A repeated musical entity**

We understand repeats easily because all characteristics of the former entity, in terms of parts and wholes, correspond to the latter entity in the same way. In information theoretic terms, the re-iterated part of a repeat presents no new information, therefore the set of instructions that explains the relationship between the two entities is minimized. Either understanding corresponds to a simple relation and the perception of continuity. Importantly, the above assumes that we perceive groupings as they are notated in Figure 3.14. If we were to group such an example differently so that we didn't perceive a repeat, then the preceding would not apply.

A musical progression is more difficult to understand, requires a longer program to represent and incurs greater cost to derive a solution when its constituent entities are more unlike, given that relationships between unlike musical entities are more difficult to describe. When the relationship between two entities is difficult to describe, it becomes difficult to create links between those entities. Links promote continuity; their absence promotes discontinuity. For example, consider the following excerpt from Coffey's *Koans* (Figure 3.15):

[Click to listen](#)

**Figure 3.15. Excerpt from Coffey's *Koans*.**

We are presented with a recording of street noise, glitch elements, duck sounds, sustained sawtooth-like hum, electronica drum beats and free-wheeling distortion. Indeed, these entities are significantly unlike, which makes the task of understanding

and describing the relationships between them difficult. In information theoretic terms, each contrasting entity presents us with significant amounts of new information, which means that describing the relationship between contrasting entities would require a lengthy list of instructions.<sup>14</sup> Using the previous definitions, Figure 3.15 thus presents us with a complex sequence, which in this case, corresponds to discontinuity.

With the above said, it is not the case that the entities that constitute difficult-to-understand relationships are inherently complex or difficult to understand themselves, rather, they become complex or difficult to understand as they form relations with other entities. Thus, the simplicity / complexity of entities is independent of the simplicity / complexity of the relations they constitute.

For example, Birtwistle's *Carmen Arcadiae Mechanicae Perpetuum* creates complex relations, and thus discontinuity, with simple entities. One such illustration of this occurs in ma. 8, where unwavering sustained notes interrupt jagged rhythmic gyrations (mm. 2-7) (see Figure 3.16).

---

<sup>14</sup> We can describe the relationship between contrasting entities in terms of *transformational distance*, which I will discuss in depth in chapter IV.

# Carmen Arcadiae Mechanicae Perpetuum

Harrison Birtwistle

Flute *ff*

Oboe *ff*

Bass Clarinet *ff*

Bassoon *ff*

Trumpet *con sord.* *ff*

Trombone *con sord.* *ff*

ophone *have sticks* *fpp* *ff* *fff*

Piano *fff*

Violin 1 *ff*

Violin 2 *ff*

Viola *ff*

Violoncello *ff*

Doublebass *ff*

Copyright 1981 by Universal Edition London Ltd. London

UE 1656 L

All rights reserved  
Printed in Austria

Figure 3.16. Harrison Birtwistle, *Carmen Arcadiae Mechanicae Perpetuum*, mm 1-8.

[Click to listen](#)

**Figure 3.17. Harrison Birtwistle, *Carmen Arcadiae Mechanicae Perpetuum*, mm 1-8.  
Perf. Alarm Will Sound**

Of the possibilities by which mm 2.-7 may be continued, the entity voiced in ma. 8 could be argued as one of the simplest, given that it comprises a small number of elements that change little in pitch, rhythm, dynamics and timbre. Yet, this relation is not particularly simple or easily understood given the context, which is characterized by contrasts in pitch, rhythm, etc. Given that simple, easily understood progression is associated with continuity, it is unsurprising that this moment is perceived as highly discontinuous.

Conversely, a complicated but repeated drum pattern, as in Figure 3.18, could characterize simple progression:



**Figure 3.18. Simple progression via complex entities.**



**Figure 3.19. Simple progression via simple entities.**

Without engaging the question of rhythmic complexity measures, most would agree that Figure 3.18 is more complicated than Figure 3.19 when considering relative number of notes, time signatures and beat divisions. Despite its relative complexity, the progression of Figure 3.18 from measure to measure is easily understood given each subsequent measure is simply a repetition of the previous measure. Figure 3.18 therefore represents simple progression (Figure 3.19 also exhibits simple progression for the same reasons).

Thus, in the case of continuity, the requirements of simple progression have more to do with the relationship between adjacent entities and less to do with an entity's inherent simplicity / complexity. We can't conceive of simplicity in an absolute sense, where fewer agents asserting fewer actions equate to acontextual simplicity. Rather, simplicity in this case describes specific kinds of change that depend on the *relative* character of entities.

Similarity is again important: the more similar the musical entities, the less change between them and the simpler the relation. In this way, musical simplicity, and the continuity that it inspires, again depends on similarity judgments. When entities are dissimilar, their relationship is likely more difficult to understand, and thus is perceived as more complex.<sup>15</sup> These complex relationships can evoke notions of discontinuity.

### 3.3.3. Regularity, Naturalness

The terms "regular" (Wertheimer: "If several possibilities are present by which a part may be continued in a whole, the simpler, more regular will be chosen" (Helson, 28)) and "natural" (Koffka: "any curve will proceed in its own natural way, a circle as a circle, an ellipse as an ellipse, and so forth" (153)) can also characterize our experience of continuity. While the words are not synonymous, they both suggest that phenomena conform to pre-existing knowledge, and thus to expectations.<sup>16</sup> Thus, for us to perceive continuity via regular or natural progression, we must already have knowledge of the entities and context that constitute that progression. The only way that we could identify a chorus following a verse in a pop song as regular or natural is if we have pre-existing knowledge of conventional pop song form. When listening to a piece in sonata form, one familiar with the conventions of 18<sup>th</sup> and 19<sup>th</sup> c. Western Art music would find

---

<sup>15</sup> This is not always the case for the juxtaposition of dissimilar entities is sometimes easily understood, particularly when such juxtapositions adhere to stylistic convention. For example, a baroque concerto typically juxtaposes a virtuosic soloist against an orchestra, which often creates moments of great contrast. One familiar with this style may not find such juxtapositions particularly discontinuous as they adhere to a particular stylistic convention.

<sup>16</sup> In the case of "regular", that conformity may be between phenomena and temporal expectations. In the case of "natural", that conformity may be between phenomena and a specific context.

a recapitulation preceded by a development section to be a regular and natural progression.<sup>17</sup> Such a listener would experience a higher-level, formal continuity in this scenario. We enhance the detail and scope of our knowledge of phenomena with experience, which affects how we will perceive that phenomena in the future.

Pre-existing knowledge can also be established dynamically in the course of a piece. If a particular passage *A* always follows a particular passage *B*, we may start to bind those two entities in a way such that *B* following *A* represents adherence to convention. For example, in John Zorn's *Road Runner*, the predominance of short, contrasting blocks establishes a convention where a short, contrasting block will be followed by another short, contrasting block (see Figure 3.10). As perceived phenomena, or some emergent characteristic of phenomena, agree with established conventions, such progressions could be considered "regular" or "natural", and thus continuous. *Road Runner* is thus continuous as its higher-level, formal organization unfolds according to dynamically-established convention, and thus is "regular", despite its lower level discontinuity.<sup>18</sup>

As notions of "regularity" and "naturalness" result from agreement between phenomenal configuration and convention, they are products of similarity judgments. The more phenomena and convention agree, the more similarity → regularity / naturalness → continuity will be perceived. In contrast with the description of simplicity above, here, entities are compared to abstract, amalgamated constructs rather than to recent phenomenal percepts. While both make use of pre-existing knowledge, judgments of "regular" and "natural" are differentiated in their dependence on such knowledge.

When juxtaposed musical entities appear as unnatural or irregular, we are potentially struck by a sense of discontinuity. Unnaturalness and irregularity occur

---

<sup>17</sup> As previously discussed, this type of agreement can potentially dilute discontinuous relations: when the progression of unlike entities becomes conventionalized, they may appear as less discontinuous (this would be interesting to test).

<sup>18</sup> Again, these are hierarchical distinctions that will be explored further in chapter IV.

when phenomena don't conform to pre-existing associational knowledge and thus, expectations. For example, imagine you are watching a performance of Beethoven's Symphony No. 6 when suddenly a performer in a purple suit playing twelve-tone passages on an electric guitar jumps out on stage, interrupting the orchestra. A variety of associations are violated in this example (musical - harmonic, formal, orchestrational; stylistic – clothing; presentational; etc.) each of which contributes to the moment's discontinuity (the strength of the discontinuity is a function of the number and type of these violations). As this contrast doesn't adhere to pre-existing associational knowledge, it contradicts our expectations (note, this would still be the case even if we knew the purple guitar player was about to emerge, given the strength of the pre-existing associational knowledge). As a result, the juxtaposition occurs to us as unnatural, irregular and thus, discontinuous.

### 3.3.4. Regularity / Naturalness and Simplicity

Regularity / naturalness are related to simplicity / complexity in that all are affected by pre-defined knowledge.<sup>19</sup> As explained above, we perceive relations to be regular and / or natural as they agree with pre-defined knowledge. In the case of simplicity / complexity, as knowledge is better defined and better understood, phenomena that relate to such knowledge can appear as simpler than if such knowledge was not well defined or well understood. Therefore, simplicity judgments can depend on pre-defined knowledge to a greater extent than previously discussed. It follows that when a relation agrees with firmly established knowledge in a clear way, we could describe it as regular, natural, *and* simple.<sup>20</sup> The extent to which that knowledge is

---

<sup>19</sup> Importantly, “naturalness”, “regularity” and “simplicity” are not synonymous, for there can be phenomena that evoke the ideas “regular” and / or “natural” that are not necessarily simple (similarly, there can be natural phenomena that are not regular and regular phenomena that are not considered natural). These examples do not contradict the preceding, rather, they exemplify different ways that the words “regular”, “natural”, and “simple” can be understood.

<sup>20</sup> It also may be the case that we can differentiate between simple regular / natural and complex regular / natural. The former would characterize organizations that agree with pre-defined knowledge in an easy to understand way; the latter would characterize agreement that is difficult to understand.



defined, which is a function of exposure and understanding, will affect perceived regularity, naturalness and simplicity / complexity.

North and Hargreaves explored the relationship between perceived complexity and pre-defined knowledge (which they describe in terms of “familiarity”) in popular music (1995). They cite Heyduk’s (1975) “optimal complexity model” that predicts that increased familiarity with a stimulus should reduce the perceived complexity of that stimulus (North; 1995, 77-78). In their own experiments involving subjective ratings for complexity, familiarity and liking of popular music excerpts, North and Hargreaves’ results support the predictions of the optimal complexity model. By including participants of various musical backgrounds, they were able to show that perceived complexity differed depending on level of musical training (1995).

These experiments support the idea that when a listener becomes more familiar with a piece or convention, he understands it with less difficulty, which can inspire a sense of continuity.<sup>21</sup> This agrees with intuition: as we become familiar with the characteristics of a particular musical style, we gain the ability to judge specific attributes as regular and natural relative to their context. When these perceived attributes match our pre-defined knowledge, we understand them without difficulty: they appear as simple. These judgments of simplicity, regularity and naturalness contribute to our sense of continuity.<sup>22</sup> Such judgments can be inspired in correspondence with a variety of musical parameters including timbre, chord progressions, harmonic content, form, orchestration, melodic structure and rhythmic configuration.

---

<sup>21</sup> One wonders about the cognitive reason for this phenomenon. Perhaps fewer resources are needed to encode the experience given knowledge of the experience already exists. Attention and reaction to surprising progressions would also differ in the case of known works potentially requiring fewer resources. The lower cognitive load could make it easier for the listener to understand a piece, which would result in the listener judging the piece as simpler.

<sup>22</sup> It might be the case that the combination of these notions creates a stronger sense of continuity than would be produced by any one notion in isolation (that is, a passage considered regular but not simple might evoke a weaker sense of continuity than a passage that is considered regular and simple).

### 3.3.5. Continuity via Directionality

Musical continuity often inspires a notion of direction, and thus can be understood in terms of dependence and implication. In the context of temporality, a necessary component of musical experience, Koffka describes continuity as a product of inter-element dependence:

A melody is a whole, organized in time. That its later members depend upon its earlier ones is the exact counterpart of the fact that the right upper quadrant of a seen circle depends on the left lower one ... The earlier notes of the melody have an effect upon the later ones, because they have started a process which demands a definite continuation. (437)

In *The Rhythmic Structure of Music*, Cooper and Meyer similarly discuss the tendency of a group's beginnings to imply a direction towards a specific conclusion (61).<sup>23</sup> Diana Deutsch writes that the Gestalt principle of Good Continuation states: "elements that follow each other in a given direction are perceived together" ("The Psychology of Music" 100-101).

These ideas - dependency, implication and direction - convey a specific type of relationship between terms where certain pairings are more strongly linked than others. That is, if *A* is more strongly linked to *B* than to *C*, then the appearance of *A* would imply or point to *B* rather than *C*. *B*, in this case, would then depend on the appearance of *A* first (implies / points looks forward, depends / follows looks back). Alternatively, we could describe this relationship in terms of probability, where the probability of *A* leading to *B* will determine the extent to which *A* implies *B*, and *B* depends on *A*. We could further illustrate this relationship through ideas such as grouping (*A* and *B* are grouped together in a specific way), tendency (*A* tends to go to *B*), expectation (when we

---

<sup>23</sup> While some of these comments are made in reference to tonal music, the concepts are still applicable to non-tonal configurations.

hear *A*, we expect *B* to follow) or even causality (*A* causes *B*).<sup>24</sup> In all cases, we are talking about a kind of directionality where an entity suggests movement to a specific place (in this metaphor, direction / movement / place could occur along a variety of musical parameters including harmony, tempo, registral destination or volume). The stronger the link between *A* and *B*, the stronger the sense of directionality towards *B* when *A* appears. When the link between *A* and *B* is sufficiently strong so that directionality is perceived, continuity is created between the two entities.

When strongly linked entities are connected through time, a path is created, which inspires our sense of directionality and continuity. As movement affirms tendencies or expectations, a directional path is created that traces past movement. As an entity in the “now” moment projects tendencies or inspires expectations about a specific future destination, we envision the path the entity must follow to reach that destination. Thus, directionality (and its associated concepts) has a past, present and future. In each of these temporal periods, a gesture is linked to a common vector or contour, which allows us to relate the constituent entities to each other. As we successfully relate constituent entities to each other relative to a common vector or contour, we can understand them as continuous.

In music, a directional path can result when the relations between individual elements are similar, as in an ascending pitch sequence. Diana Deutsche points out that adjacent pitch elements that occur along a unidirectional pitch vector will tend to be grouped together in accordance with the principle of continuity (1986, 32-39). If the interval between each pitch is identical, as in Figure 3.20, continuity is perceived because of relational sameness.

---

<sup>24</sup> In regard to causality, one could describe these relationships in such terms, but it would be a weak kind of causality if one at all. If one were to argue this point, he would say the I chord in a V – I cadence occurs because the V chord occurred first, or, the V chord causes the I chord to occur. That is not to say that the I chord *must* appear as a result of the appearance of the V chord, as is common in cases of physical causality, but there is a strong tendency for V to follow I (because of stylistic convention), so when we reflect on the V - I cadence after the fact, we can describe I as occurring because of the occurrence of V.



[Click to listen](#)

**Figure 3.20. Pitch Sequence - identical intervals.**

If the intervals between pitches are not the same, as in Figure 3.21, an abstracted “ascension” can still characterize each relational pair. This ascension creates a sense of similarity between pairs, which is sufficient to create a sense of continuity between individual elements and the sequence as a whole.



[Click to listen](#)

**Figure 3.21. Pitch sequence – non-identical intervals.**

In both cases, “ascending” is not a quality of any particular pitch, rather, it is a quality of the relationship between particular pitches. Our sense of ascending direction inspires our sense of continuity, which is attributable to a specific kind of relationship between phenomenal events (similarity), *rather than* to the phenomenal events themselves. In this way, continuity exists on an abstract level, which can be described in terms of a hierarchy (I will explore this idea further in chapter V).

Of course, this type of relational continuity is not limited to ascending pitch sequences. We can perceive downward motion and, further, more complex motions that remain continuous despite moving in multiple directions over time (changing contours). We can also apply motions to other musical parameters: a vector could represent a louder volume, or a faster tempo, for example. Directions in different dimensions are also possible, depending on the musical parameter under consideration.

### 3.3.5.1. *Directionality and Schema*

We can also perceive directionality if we are fluent with the conventions of a particular style or piece. We gain such fluency by becoming familiar with frequency distributions. For example, as we become familiar with Western tonal music, we gain a sense of the kinds of chords that most often follow a V7 chord and the kinds that do not. Such knowledge allows us to describe such conventions in terms of tendencies: a V7 chord tends to go to I. A listener fluent with these conventions and tendencies generates expectations about the direction of a sequence of events before that direction has come to fruition. James Carlsen investigated how conventional knowledge affects directional expectations by playing a series of melodic fragments to participants who were music majors in professional music schools in Hungary, Germany and the USA, and then asked them to sing “what they believed would have been the continuation of that melody had it not been interrupted” (15). Results showed that “cultural milieu has considerable influence upon melodic expectancy” (Carlsen, 20). It appears that expectancies, and therefore perception of directionality, depend on a specific type of contextual knowledge that is informed by individual and cultural experience.

David Huron describes this type of knowledge in terms of mental schemas, which are generalizations formed from exposure to a number of individual instances (225). These schemas are invoked by musical cues, such as instrumentation, which guide our expectations in the course of listening (Huron, 204). As these schema-driven expectations are met as the music unfolds, direction and continuity are created. For example, within certain genres of rock music, relatively subdued verses (defined by clean guitars, moderate dynamics and relatively reserved vocal articulations) are often juxtaposed with intense and energetic choruses (defined by distorted guitars, fortification via additional instruments, loud dynamics and vocal straining). There is usually no transition between these sections, so one could imagine that those unfamiliar

with such music would find such a juxtaposition to be striking and discontinuous. For those familiar with these conventions, such juxtapositions are fully regular and natural and thus, continuity is preserved despite significant contrast. It follows that violations of these schema-driven expectations can create discontinuity. These sorts of (dis)agreements between phenomena and established knowledge are similar to the kind described in the “regularity” and “naturalness” section. Therefore, here, as was the case for “regularity” and “naturalness”, similarity judgments are keenly important to the perception of continuity.

### 3.3.5.2. Directionality via Dynamic Expectation Creation

Conventions can also be created dynamically in the course of a piece through sufficient repetition, which influences future expectations and our sense of (dis)continuity.<sup>25</sup> When the repetition of a gesture or configuration results in a convention, the partial statement of that convention implies a particular direction (we, as listeners expect that certain entities will follow specific antecedents) and the full statement of that convention leaves a trace of that direction. For example, in John Zorn’s *Road Runner* (Figure 3.10) and Naked City’s *Snagglepuss* (Figure 3.11) the repeated presentation of short, contrasting blocks leads us to assume that the current short, contrasting block will be followed by another short, contrasting block. When we indeed hear another short, contrasting block, our expectations are confirmed, which gives us a sense of a “good fit” between what actually occurred and what was implied. This agreement can lead to a notion directionality and continuity.<sup>26</sup>

In *Sweet Anticipation*, David Huron discusses the relationship between exposure (which is related to repetition) and expectation in the context of cross-cultural psychological studies (one example of the latter is Kessler, Hansen & Shepard (1984),

---

<sup>25</sup> There may be ways other than repetition to establish such generalizations, such as through accent.

<sup>26</sup> In the case of both musical examples, this “higher-level” continuity is at odds with the discontinuity that exists between each musical block. I will discuss these kinds of hierarchical relationships in greater detail in chapter V.

who used the probe tone method to test melodic “goodness of fit” by playing both Balinese and American participants melodies based on Balinese and Western scales) (168-172). Huron concludes that naïve listeners are sensitive to frequency of occurrence in a stimulus while enculturated listeners invoke appropriate schema and evaluate fitness on the basis of “frequency of occurrence in the totality of their past listening exposure” (170). In both cases, expectations and the perception of fitness are affected by frequency of exposure. The former case is particularly supportive of the notion that within a musical piece, repetition can dynamically affect the perception of future musical events, thus creating a sense of directionality. This idea is particularly relevant in the case of contemporary music, where pieces typically avoid conventional forms and gestures.

The idea of dynamic expectation creation has also been explored in the context of information theory. In their article “Information as a Measure of the Experience of Music”, Kraehenbuehl and Coons propose a model where the overall probability system that guides our expectations is affected by each experienced event (3). When we encounter musical events, our expectational (or probability) system is either confirmed or nonconfirmed (to use Kraehenbuehl and Coons’ language), which produces various amounts of information and affects future expectations. Confirmation / nonconfirmation of these expectations creates perceptions of continuity (and directionality) and discontinuity, respectively.

The psychological studies cited by Huron and the model proposed by Kraehenbuehl and Coons highlight how our expectations are informed not only by schema, but also by how often entities occur and how they are configured in a specific piece. If entities are repeated enough in the right configuration, listeners’ expectations of the future and sense of directionality will change. As those expectations are confirmed or thwarted, continuity or discontinuity can be created. Composers can thus use repetition to affect expectations and create (dis)continuity dynamically in the course of a

piece. This is particularly important to contemporary music that doesn't evoke specific schema, or seeks to create new types of relationships.

### Summary

Musical discontinuity is typically defined in terms of breaks / gaps (which can be created by tacets, juxtaposed contrast, surprise and change) and the negation or lack of continuity. Continuity, then, describes an entity whose parts are connected without interruption or breaks. We can perceive such entities as continuous, as having no breaks or interstices, when a notion of sameness, or sufficient similarity, characterizes our perception of them at *every* point along a dimension relative to other points (*uninterrupted continuity*). These conclusions of "sameness" result from judgments of difference. Because we are able to perceive difference within unbroken wholes, cases of judged identity are better described as similarity comparisons that satisfy the appropriate property-sharing requirements. When (dis)similarity comparisons do not satisfy the appropriate property-sharing requirements, we perceive discontinuity. It follows that (dis)similarity comparisons significantly influence whether we perceive continuity or discontinuity within whole entities.

Continuity can also exist between entities that we perceive as distinct. We can perceive continuity between distinct entities when we make connections between them, which we do by bringing them into a relation. We can then describe the structure of continuity relations as including two distinct entities that are separated by an *intermediary*. The content of a relation's intermediary affects how we perceive (dis)continuity. Intermediaries that contain silence can produce salient discontinuities. When intermediaries contain *sound*, we must differentiate between the kinds of content that make us perceive a transition between *A* and *B* versus the kinds that inspire us to perceive a separate entity *C*. Psychological research that investigates how intermediaries affect continuity perception may also be of use in these descriptions.



One important kind of continuity that involves distinct entities and a contrasting intermediary is *interrupted continuity*, which we can represent as (*entity A* → *entity B* → *entity A*). In the field of auditory perception, interrupted continuity has been studied in the context of the “continuity illusion”, which essentially involves a softer sound that is interrupted by a louder sound (often a noise burst), which we can represent as ABA. The continuity illusion reinforces the idea that the content of intermediaries plays an important role in (dis)continuity perception. In the wider world of music, interrupted continuity has been explored in a number of contexts, including audio production techniques and compositions such as Ives’ polytextural *Central Park in the Dark*.

When two entities form a relation, their temporal position also has an important influence on (dis)continuity judgments. Specifically, discontinuity percepts most often seem to occur when entities are separated by short intermediaries. Conversely, short intermediaries can also bring contrasting entities into a relation that promotes a sense of continuity. Longer transitions that connect the entities they mediate in a linear way typically promote continuity. In sum, the temporal characteristics of intermediaries, in combination with other elements such as the relationship between contrasting entities, higher-level organizations and context, affect how we perceive (dis)continuity in music. Discontinuity can be also produced by distinct entities that have some overlapping simultaneity.

We can also describe the characteristics of continuity in relations in holistic ways, which include sameness, simplicity, regularity / naturalness and directionality. In each of these cases, the perception of similarity seems to have significant influence on the perception of continuity. When we perceive similarity between components of a whole, distinct entities, phenomena and pre-defined knowledge (which include schemas and dynamically-created expectations) and relations on abstracted levels (directionality), notions of continuity can be inspired. It follows that understanding the relative similarity between terms will be important if we are to understand how those terms

create notions of (dis)continuity. In all cases where the requirements of continuity are specified, discontinuity can be created by interrupting those requirements.

## 4. Chapter IV: Perceiving musical (dis)similarity

*Similarity is at the very heart of music, not only because it allows connections and associations to be made between different musical entities, but because it enables the emergence of the core musical entities themselves* (Cambouropoulos, 2009, 8).

Our (dis)similarity percepts significantly influence our experience of musical (dis)continuity and music in general. Irene Deliege, an influential writer on the subject, describes the importance of similarity perception as it can “give the listener the feeling that he understands what he is listening to without having to study the compositional rules on which the music was based” (2007, 13). Ultimately, this understanding inspires notions of coherence and unity that are attributed to the piece.<sup>27</sup> Cambouropoulos (2009) highlights how (dis)similarity percepts allow us to break a sound stream into parts and make associations between those parts, which inspire larger perceptual entities and structures. Tversky (1977) reflects on similarity’s foundational role in perception: “It serves as an organizing principle by which individuals classify objects, form concepts and make generalizations” (327).

It follows that we need to understand why we perceive sounds as (dis)similar if we are to understand why we perceive them as (dis)continuous. The key to this understanding will be the relationship between perceived likeness and difference. The Gestalt law of similarity states that “configurations are formed out of like elements” (Deutsche 1982, 100). As previously stated, conclusions of “likeness” result from judgments of difference: when difference is detected, we are able to draw a boundary demarcating two distinct entities, which we can then compare to determine the entities’

---

<sup>27</sup> Deliege, via Meyer, discusses the relationship between coherence, unity and similarity coming to fruition as a function of Romantic ideals and novice audiences.

relative likeness.<sup>28</sup> Similarity thus depends on the perception of both likeness *and* difference. Deliege (2007) reflects this quality in her definition of similarity: “a process that compares two or several things ... This process is required to assess ‘resemblance’; in other words, it must reveal those characteristics by which the objects that are being compared differ ... or it must show that there is pure and simple ‘identity’ or a state of maximum similarity” (9-10). Similarly, Cambouropoulos (2009) states that “similarity is often defined as partial identity, that is two entities are similar if they share *some* properties, but not necessarily all” (9). In summary, similarity involves two (or more) entities and a relation between them that illuminates likeness and difference.

Articulating similarity’s constituent likeness and difference is often difficult, even though we make musical similarity judgments on a regular basis. When listening to music we are often enlightened by the notion that the current theme is similar to the one first heard in the piece; that the current band sounds like another band; that the current melody is reminiscent of one heard before. Given the immediacy and tangibility of such judgments, we could be forgiven for thinking that the relations underlying those judgments could be easily described. In practice though, similarity proves to be a complex subject, which makes the process of describing absolute “laws” that govern similarity relations a difficult task. How similar is a clarinet sounding A4 to a violin sounding B7? How similar is the sound of a piano to the sound of a lute? How similar is a major chord played by an electric guitar to a salsa rhythm played on hand drums? These are questions whose answers are typically characterized by nebulous intuitions and not specific descriptions.

---

<sup>28</sup> This relationship raises complicated causal questions. Is it the case that a sense of difference causes us to segment a stream into two distinct entities after which we can then compare those entities to gauge relative (dis)similarity? The problem with such an account is that the “sense of difference” is itself a (dis)similarity judgment, that is, we distinguish entities by judging their relative (dis)similarity. We conclude they are distinct when they exhibit some required degree of dissimilarity, which we typically understand clearly and immediately. How is this possible? How can we compare two entities to gauge their relative (dis)similarity when the two entities haven’t yet been defined? We can also describe this scenario as “separate” but related processes: in one way, we clearly and immediately understand the relative (dis)similarity between distinct entities, thus their distinction. In another way, we don’t immediately understand how entities are related beyond their distinctness, which means we need to reference a richer network of information to enact a (dis)similarity comparison. While such questions don’t necessarily have to be answered for us to consider (dis)similarity perception more deeply, they do indicate the complexity of the topic at hand.

Such difficulties have inspired criticism of similarity as an explanatory construct. Perhaps the most famous of these was made by philosopher Nelson Goodman (1972). He points out that similarity is only meaningful when one can describe “in what respects” A is similar to B. That is, the statement that song A is similar to song B is devoid of meaning until we specify *how* A is similar to B (perhaps they have similar chord progressions, tempo, or rhythmic configurations). In each case, meaning is supplied by the discovery of the particular relevant properties and relations, not by the abstract container “similarity”. If we are to jettison similarity as a general construct and instead focus on specific relations between properties, we are still confronted by difficulties. It seems that similarity perception depends on a number of factors including context (social and musical), priming, memory, temporal perception (Wiggins, 2007), attention, individual differences, individual variation (opinions can change over time or vary with mood and other factors) (Berenzweig et al., 2004), and hierarchical representation. Because of this, a variety of judgments are possible from the same compared phenomena, making the idea of similarity perhaps untenably flexible. Researchers have pointed out how this flexibility prevents similarity from being a useful theoretical construct to provide the foundation for other cognitive processes such as categorization (see Medin et al., 1993 for further discussion).

Nonetheless, our sense of similarity as a foundational part of musical experience persists, and as such, fuels our indefatigable drive to explore the subject further. Instead of being defeated, we can use the aforementioned critiques as a guide to avoid or address the problematic parts of similarity as a construct. Indeed, researchers and thinkers have adopted this perspective; the following explores some of what they have to say:

## 4.1. Models of Similarity

#### 4.1.1. Spatial Model

In the spatial, or geometric, model of similarity, exemplified by multidimensional scaling (Shepard, 1957, 1962a, 1962b), each entity is represented by a point in a multidimensional feature space. One can determine the degree of similarity between entities by measuring the distance between the representational points. Spatial models assume minimality [ $D(A,B) \geq D(A,A) = 0$ ], symmetry [ $D(A,B) = D(B,A)$ ], and the triangle inequality [ $D(A,B) + D(B,C) \geq D(A,C)$ ] (Guan & Wang, 2008; Tversky, 1977). Such an approach is attractive in that it allows us to represent and quantify similarity relations along a number of dimensions (dimensional representation will be explored later in this chapter and also in chapter VI). The problems of spatial models involve the aforementioned assumptions, which are addressed by featural models.

#### 4.1.2. Featural Model

Psychologist Amos Tversky's featural model challenges the core assumptions of minimality, symmetry and the triangle inequality of geometric models (1977). Minimality cannot be assumed given the possibility that different kinds of identical entities may result different distributions of same / different judgments. This is musically relevant: recognizing identity between a theme and its re-statement depends on the character of the theme itself (such as complexity, duration and memorability) as well as the relationship between the two entities (such as the intervening duration between the two and contextual surroundings of each). Symmetry can also not be assumed in similarity judgments, for we (in contemporary Western culture) say that some nebulous sound is similar to a piano, not that a piano is similar to some nebulous sound. Tversky points out "this asymmetry in the *choice* of similarity statements is associated with asymmetry in *judgments* of similarity" (1977, 328). The triangle inequality is also problematic as it implies that if A is similar to B and B to C, then A

should be similar to C. We can see how this is problematic in music: if a violin is similar to a guitar because they are both string instruments, and a guitar is similar to a drum machine because they are both regularly used in rock contexts, that is not to say that a violin is similar to a drum machine.

Instead, Tversky suggests that similarity between two entities is a function of their common and distinctive features. Features can correspond to individual components, concrete properties (size, color), or abstract attributes (quality, complexity). For two entities *A* and *B*, similarity is a function of the features that are common to *A* and *B*, the features that belong to *A* and not *B*, and the features that belong to *B* and not *A* (See 1977 for a more comprehensive description of the model).

Featural models also have problems, some of which are discussed by Emiliós Cambouropoulos (2001, 2009) and Dedre Gentner (1983). Cambouropoulos points out that entities require elaborate descriptions, that is, “that the individual sets of all the features that are important for the description of each object be precisely defined (rather than using only one general set of features that accounts for all the objects)” (2001, 351-352). In addition, the point at which the identified feature set is sufficient to make comparisons is not specified. It seems that if one tried to describe the entire feature set of any entity, the list would be exceedingly long if not infinite. Cambouropoulos also points out that it is not clear how one determines which properties are relevant in a similarity comparison (2001, 351-352). Relatedly, Gentner (1983) shows how in the case of analogy, the most salient features of the comparison, and not the agreement of the entire featural set, are the only ones that really matter. She gives the example of the statement “a battery is like a reservoir”. The essence of this comparison is that both store potential energy and subsequently make energy available. The size, shape, color or other features that may or may not correspond between the two is of much less importance (Gentner, 1983, 156).

Additionally, both spatial and featural models do not capture the structure of

representations. That is, in the case of a red square above a blue triangle, one must bind “red” to “square” and “blue” to “triangle” and also must orient each properly in regard to the “above” relation. Spatial and featural models have difficulty making these connections both between properties and entities and between entities and relations (see Larkey & Markman, 2005). One can imagine how such representations are important in music. If we have a distorted guitar playing C4 and a “clean” violin playing G4, we could describe the relation in terms of orchestration (guitar / violin), level of distortion (distorted / clean) and pitch interval (5<sup>th</sup>). Such a featural description does not capture the fact that the guitar is distorted, or that the “clean” sound is intertwined with the pitch G4. Such connections require structured representations.

#### 4.1.3. Structured Representations

Structured representations address some of the deficiencies of spatial and featural models by illuminating how the attributes and relations of *entity A* correspond or align to the attributes and relations of *entity B*. Two approaches that address structured representations are *Structure-Mapping* and *Transformational* models.

##### 4.1.3.1. Structure-Mapping Models

Structure-mapping models, as described by Gentner (1983) in the context of analogies, illuminate alignment and correspondences<sup>29</sup> by mapping the *attributes* and *relations* of *A (base)* to *B (target)*, which highlights commonalities and differences. Given the sufficient conditions, we can then say: “an *A* is like a *B*”. Similarity, the present topic, “is like analogy, in that both involve an alignment of relational structure” (Gentner & Markman, 1997, p. 48; Gentner & Markman, 1995). The difference between the two involves how relations and attributes are mapped, which result in the following

---

<sup>29</sup> A *correspondence*, as defined by Goldstone (1994), is “a connection or association between parts from two scenes” (3).



categories as defined by Gentner (1983): *literal similarities* (a large number of entity attributes and relations are mapped between base and target relative to nonmapped predicates), *analogies* (relations but few or no entity attributes are mapped), *abstractions* (the base is an abstract relational structure; there are no nonmapped predicates), *appearance matches* (entity attributes but not relations are mapped), and *anomalies* (neither entity attributes nor relations are mapped). Figure 4.1 represents how these various categories relate to each other.

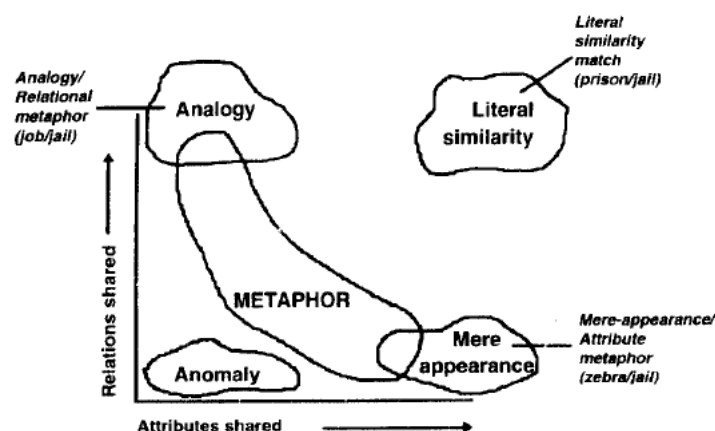


Figure 4.1. "Similarity space, Showing Different Kinds of Matches in Terms of the Degree of Relational Versus Object-Description Overlap" (Gentner & Markman, 1997, p.48)

When considering the alignment of relational structures, we must contemplate the question of which predicates are mapped from base to target. Gentner answers this question with the *systematicity principle*, which states that a predicate that is part of a deeply imbedded, interconnected system is more likely to be mapped than a predicate that is not (see Gentner, 1983 for further explication). We can use such a principle to guide our focus towards relevant attributes and relations.

With relational structures aligned, we are able to talk about inter-entity commonalities. Commonalities that form a relational system are more important than ones that do not. For example, imagine two pictures: in picture A, a man lifts a box

while in picture B, a man lifts a weight. In both examples, a relational system formed by commonalities is established, in this case a man lifting an object. A commonality that is not connected to the relational system is exemplified by a chair in the right background of picture A and a chair was in the left foreground of picture B. We can also distinguish between a match between corresponding or aligned elements, called a match in place (MIP), and a match between elements that don't correspond or are not aligned, called a match out of place (MOP). For example, imagine a cube A, which has a blue top side and a green right side, and cube B, which has a blue top side and a green front side. The blue top would be MIP (blue match on corresponding top sides) while the green sides would be a MOP (green match on sides that don't correspond). MIPs influence similarity perception more than MOPs (Goldstone, 1994; Larkey & Markman, 2005).

Structural alignment similarly affects the salience of perceived difference. When two different entities play the same role in a matching structure, they can be described in terms of *alignable differences*. In the previous example, the object being lifted (part the relation that is shared between both pictures) is characterized by two different entities: a box in one and a weight in the other. *Nonalignable differences* refer to aspects of one of the compared terms that do not correspond to any aspect of the other compared term (for example if there was a mural on the wall of picture A but nothing on the wall of picture B). As in the case of commonalities, binding entities to a relational system makes alignable differences more important than nonalignable differences. In support of this statement, it has been experimentally shown that people list *more* differences for similar pairs than for dissimilar pairs, and they indicate more alignable differences for similar pairs than for dissimilar pairs (this is counterintuitive). It is thus suggested that alignable differences impact similarity perception more than nonalignable differences (Gentner & Markman, 1997). If we are to capture the essence of similarity relations then, we must consider structural alignment, commonality and difference between compared entities.

Various models of this structure mapping process have been proposed and include the *Structure Mapping Engine* (SME), *Similarity as Interactive Activation and Mapping* (SIAM), *Connectionist Analogy Builder* (CAB) and *Systematic Similarity Model* (SSM) (see Larkey & Markman, 2005; Guan & Wang, 2008 for more information and experiments that compare these models). As these models typically have focused on linguistic and / or pictorial inputs rather than musical ones, I will not explore them in great depth here. With that said, many of the preceding ideas are relevant to music, after all, music is also constituted by attributes, relations and structures. I will therefore use the conceptual parts of these theories and models that help us describe musical (dis)similarity and (dis)continuity).

We can apply the preceding ideas to musical contexts in a variety of ways. Musical similarity relations, like linguistic analogies, also involve mapping from a base to a target. “Abba’s *Waterloo* [target] sounds like the Foundations’ *Build me up Buttercup* [base]” (see [soundsjustlike.com](http://soundsjustlike.com) for other intriguing musical analogies). Music can also be described in terms of *attributes* and *relations*, so the idea of *structural alignment* is also relevant. For example if we were to compare Sony Rollins’ version of *My Reverie* (Figure 4.2) and Duran Duran’s *Ordinary World* (Figure 4.3), we could identify an entity, characterized by sharp transients and short decays, that articulates meter, tempo and rhythmic accents (drums); an entity that provides a rhythmic and harmonic foundation (bass); an entity(ies) that provides harmonic / rhythmic detail (piano and guitars, respectively); and an entity that provides central, prominent melodic phrases (saxophone, voice). In each example, each of these parts stands in specific relation to the others, creating a distinct structure. The two examples given have similar relational structures and thus can be aligned in a similarity comparison.

[Click to listen](#)

**Figure 4.2. Sony Rollins Quartet – excerpt from *My Reverie*. Composer: Larry Clinton.**

[Click to listen](#)

**Figure 4.3. Duran Duran – excerpt from *Ordinary World*.**

We thus can apply Gentner's structural alignment categories to musical situations. *Literal similarity*, which involves mapping of attributes and relations, is clearly musically meaningful. For example, imagine *My Reverie* is juxtaposed with *Paul's Pal* (both tracks are contained on *Tenor Madness*). The two entities share many attributes (genre, performers, instrumentation, rhythmic configurations, etc.) and relations (a solo instrument states the theme; the piano provides harmonic accompaniment relative to the soloist, etc.). Given the large number of predicates that can be mapped, we can consider this an example of literal similarity.

*Analogies*, where relations are mapped but few or no attributes are mapped, can also describe the way in which musical entities align. Consider *My Reverie* and *Ordinary World*. Both pieces of music share a similar relational structure: a solo instrument states the theme and commands most of the attention, chord-playing instruments provide harmonic structure, percussion instruments articulate beat and meter, the bass provides a low-register foundation both rhythmically and harmonically. At the same time, the attributes of each differ in terms of pitch, duration, tempo and harmony.

*Appearance matches*, where entity attributes but not relations are mapped, also seem relevant in music. For example, consider Vitamin String Quartet's version of Radiohead's *Palo Alto* versus Ligeti's *String Quartet No. 2*. The two examples share some attributes, such as instrumentation, but the relations between the instrumental voices in each piece differ significantly across a broad set of categories including pitch material, rhythmic organization, dynamic contour and timbral articulation.

*Anomalies*, where neither attributes nor relations are mapped, are clearly in play. The first 45 seconds of Coffey's *Koans* explores a variety of structural relations, one of them being anomalies. The latter occur in a variety of places, such as at 0:21. Here a

larger group that is characterized by rapid, ephemeral, and volatile gestures is juxtaposed with a single static, pure high-frequency tone (see Figure 4.4). Other than both inhabiting a very general electronic sound world, the two moments share few attributes and few relations.

[Click to listen](#)

**Figure 4.4. Excerpt from Coffey's *Koans*.**

The *systematicity principle* is also relevant given deeply embedded, interconnected networks of knowledge affect musical judgments. Such a principle could help us understand why we privilege some predicates over others, and how context plays a role in that process. For example, if a listener has expertise in piano timbres, he may privilege a timbrally-based predicate over a predicate based on another kind of parameter, such as pitch or rhythm. This choice, or tendency, could be explained by a version of the systematicity principle, which states that because this listener has a deeply embedded, interconnected base of knowledge on the subject of timbre, timbrally-based predicates are more salient in (dis)similarity relations than predicates characterized by other parameters, such as pitch or rhythm, for which the listener has less deeply embedded, interconnected bases of knowledge. Thus, the systematicity principle can become part of our contextual considerations.

Because we can use structural alignment to describe musical relations, we can illuminate commonalities and differences as previously described. In terms of commonalities, the ideas of MIP and MOP can be applied to music. Distorting the saxophone in *My Reverie* and the vocal in *Ordinary World* would create an MIP, while distorting the bass in *My Reverie* and the guitars in *Ordinary World* would create an MOP. In both of these examples, the *matches* reflect attribute similarity while the *place* (in or out) reflects the position of those attributes within a relational structure. We can further distinguish commonalities that aren't connected to relational structures. For

example, imagine we add the following drum part to both songs at a tempo that differs from that of both of them (Figure 4.5):



**Figure 4.5.** Commonality that doesn't correspond to the relational structures of *My Reverie* or *Ordinary World*.

The presence of the same part in both pieces creates a commonality between them. At the same time, the rhythmic characteristics of Figure 4.5 differ significantly from the relational rhythmic structure within each song in terms of tempo, meter and rhythmic configuration. We can then speak of three kinds of commonalities: matches between attributes in the same relational position, matches between attributes in different relational positions, and matches between attributes in positions that don't correspond to relational structures. In analogy, commonalities that form a relational structure are more important than those that do not, and MIPs are more important than MOPs. That may or may not be the case in musical relations. For example, it may be the case that Figure 4.5 as applied above influences (dis)similarity judgments more than commonalities that correspond to relational structures. Such conclusions would seem to depend on the characteristics of the entities in question.

In terms of differences, we could describe the saxophone in *My Reverie* and the voice in *Ordinary World* as alignable differences (both serve a common relational function but are sounded by different entities). We could describe the "walking" attribute of the bass part in *My Reverie* (a higher-level emergent attribute created from specific pitch and rhythmic relationships) as a nonalignable difference, as there is nothing in *Ordinary World* that serves a similar function.

Structure mapping models are useful as they give us a framework to compare

attributes and relations within and between entities. This allows us to describe a relation in terms of the roles that its components play relative to other components. This is useful to us as in music, instruments and voices play particular roles relative to other instruments and voices to articulate form, define rhythmic structure, establish harmonic territory, project melodies, etc. We can further enhance our descriptions of musical roles and structural positions by locating them in the arena of perceptual hierarchies, which I will discuss in chapter V. These various concepts will help illuminate the richness of musical relations, which will allow us to better understand how we make (dis)similarity and (dis)continuity judgments.

#### **4.1.3.2. Transformational Models**

Transformational models, as described by Hahn et al. (2003), use the *Representational Distortion* (RD) approach to similarity perception. “According to RD, the similarity between two entities is a function of the ‘complexity’ required to ‘distort’ or ‘transform’ the representation of one into the representation of the other. The simpler the transformation distorting one representation to the other, the more similar they are assumed to be” (2). They draw on Kolmogorov complexity theory to describe the complexity of representations and transformations between them. “In intuitive terms, according to Kolmogorov complexity theory, the complexity of a representation,  $x$ , is the length of the shortest computer program that can generate that representation” (2). In this context, we can take a program to mean a list of instructions. It follows that short programs correspond to simple representations / transformations and longer programs correspond to more complex representations / transformations. For similarity relations, simple transformations (short programs) equate to more similar entities while complex transformations (longer programs) equate to dissimilar entities (also see Orpen & Huron, 1992, for examples of computing the transformational distance between musical entities).

Transformational approaches have several advantages over featural and spatial accounts. Transformational models allow us to deal with structured representations (which involve attributes and relations), which spatial and featural models have difficulty with (see Hahn et al. 2003). Another important difference between transformational models and featural models is that transformations are *relations* between entities, not properties of them (features are properties of an entity). Given that similarity can be described in terms of transformations and that transformations cannot be reduced to features, this observation gives a transformational approach more descriptive and predictive power than a featural account (Hahn et al., 2003, 25).

Another advantage, which Hahn et al. (2003) point out, is that transformational models are sensitive to the fact that features sometimes become perceptually salient as a result of a transformation. For example, imagine you see a circle in a square, after which you are shown a blue circle in a square: the color of the first circle, white, is now salient as a result of the transformation (16). It follows that features are influenced by transformations, which is an important distinction versus the model that considers features wholly distinct and unconnected from their contextual surroundings. In this model, entity representation and transformational relations have a symbiotic relationship: transformations affect the parts of an entity that are perceptually salient and the parts of an entity that are perceptually salient affect the kinds of transformations that can occur (20).

Hahn et al. (2003) do not view spatial and featural accounts as “wrong”, rather, they see them as not flexible and comprehensive enough to account for the varied facets of similarity. In fact, elements of these models such as changes to a feature, feature insertion / deletion or changes along a continuous dimension have descriptive and predictive power, but such elements can be described as transformations. Further, there are many possible transformations that featural and spatial models cannot capture (25-26).



The idea of transformation is applied widely in compositional and analytical practice. Classical theme and variation form is built on the idea of starting with an established entity, the theme, and transforming it and its surrounding context in a variety of ways over the course of a piece. Louis Andriessen's *Zilver* (Figure 4.6), from the highest level of description, is a transformation from sparse to dense over the duration of the work. Nancarrow's *Study No. 21* (Figure 4.7) explores simultaneous but independent rate transformations in two voices.

[Click to listen](#)

**Figure 4.6. Andriessen's *Zilver*. Perf. The California E.A.R. Unit.**

[Click to listen](#)

**Figure 4.7. Nancarrow's *Study No. 21*.**

In terms of analytical practice, David Lewin, in *Generalized Musical Intervals and Transformations* (1987), developed a mathematically-based music theoretic approach that explored some of the possibilities of musical transformations, focusing on the relation between musical objects rather than on just the objects themselves. We can apply Lewin's transformational approach to the kinds of relations that we are interested in here, but in doing so, we must be mindful that the transformations that are salient when looking at a score are not necessarily the same as when listening to the music.

Because we can apply the idea of transformations to a wide variety of musical parameters in a wide variety of musical styles, the transformational model is a powerful tool with which we can consider musical similarity relations. Even if we don't strictly adhere to the mathematics and algorithms of Kolmogorov complexity theory, we can still use the general idea of transformational distance to describe the relationship between musical entities. Whatever approach we take, we will likely describe the transformation between whole entities in terms of their constituent parts, which we can

describe in terms of musical parameters (the focus of Chapter VI).

#### 4.1.4. Synthesis – Structural Transformation

Because both the structure mapping and transformational approaches offer useful tools for considering musical similarity (and continuity) relations, we can utilize either in our compositional and analytical efforts, depending on what is contextually appropriate. Structure mapping models give us a way to highlight structural function within and between entities, particularly when the attributes and relations of compared entities can be aligned. If we are interested in a musical context whose components align, then structure mapping can reveal much about the relative similarity of those components. Transformational models give us a way to qualitatively and quantitatively describe the relationships between compared entities in a way that structure mapping models do not, which is useful when attributes or relations do not align or correspond. If we wish to create or describe such relations, as composers and analysts of discontinuous music often do, then a transformational approach will illuminate relative (dis)similarity.

We can also combine these approaches by describing the distinctions of structure mapping models in terms of transformations, to which we can then assign weights depending on perceptual salience (Hahn et al., 2003 also discuss the need for weighting of transformations). For example, the transformational distance between a MIP may have a different weight than that between a MOP (a quantity that is captured neither by a transformational account that is not informed by structure mapping theory nor by structure mapping theory not informed by a transformational account). In the previous examples of *My Reverie* and *Ordinary World*, consider a chorus effect added to the saxophone and the voice, creating a MIP. Now consider another scenario where the same chorus effect is added to the piano in *My Reverie* and to the guitar in *Ordinary World*, creating a MOP. If we accept that MIPs influence similarity judgments more than

MOPs, then transformational distance for the MIP example would be weighted higher than the MOP example, even though the non-contextual transformations are the same (adding a chorus effect to a voice / instrument). We can describe and weight various kinds of transformations according to the relevant parts of structure mapping theory earlier defined, so that hierarchical position, structural alignment, literal similarities, analogies, appearance matches, anomalies, the systematicity principle, relational and non-relational commonality, MIP / MOP, and non-alignable / alignable differences all affect transformational weights. Importantly, we need to distinguish between transformations that are perceptually relevant and those that are not, even if the latter are capable of clear descriptions.

The models, in isolation or combination, provide complex but powerful tools to describe and predict musical similarity judgments. They give us a way to talk about entities, attributes and relations, via structure mapping and transformation. Such tools only allow us to paint part of the picture though. With them, we can depict important characteristics of a similarity relation, but within a somewhat narrow and isolated frame. To complete the larger picture, we need to describe the other factors that influence our perception of similarity, which involve musical and psychological context. When we are sensitive to context, we can choose the analytical or compositional approach that is most appropriate given the musical environment.

## 4.2. (Dis)similarity Perception and context

When describe musical (dis)similarities, we must consider not only the qualities found *within* compared entities, but also the musical and psychological context in which compared entities occur. While analyzing entities out of context may reveal something about identities and relations, it doesn't illuminate the whole of our musical experience: real music is not excerpts in isolation. Instead, real music is a stream of sonic information where the textures and ideas presented in the past influence how we

perceive and interpret textures and ideas in the present. It follows that we must be sensitive to entities, attributes, relations *and* the context in which we perceive these elements, given the music surrounding two entities may influence the perception of their (dis)similarity as much as qualities found within those entities. We can ask similar questions about psychological context, about how our tangential perceptual and cognitive processes, such as attention and categorization, affect our perception of (dis)similarity. While these are admittedly big questions that require more comprehensive consideration than is possible here, I will approach some of these topics in order to highlight the various factors, outside the qualities of the compared entities themselves, that affect (dis)similarity perception.

#### **4.2.1. Psychological Context**

##### **4.2.1.1. *Attention***

In regard to psychological context, attention plays an important role in how we perceive musical (dis)similarity. (Dis)similarity percepts can be active in guiding our attention to a particular entity or relation, and they can also result from our attention zooming in on a part of a musical texture. By considering how we attend as listeners and how entities become the objects of our attention, we will gain insight into how (dis)similarity percepts function.<sup>30</sup> Attention in listening involves both active and reactive processes. We can actively direct our attention towards a component or quality of a piece: sometimes we choose to focus on a particular instrument, sometimes we focus on emergent entities such as melodies and rhythms. We also have listening modes where we are reactive, evidenced when a particular element sticks out from a musical surface and grabs our attention, compelling us to focus our attention on that element. Of course, the actuality of the listening process is not one where our active and reactive

---

<sup>30</sup> This is particularly true if we want to take an computational approach to (dis)similarity analysis.

capabilities are mutually exclusive. Rather, over the course of a piece, we are both active and reactive listeners such that phenomenal musical attributes *and* our individual perceptual sensitivities play a role in what musical components *become* salient in our attention. This question of how and why certain musical entities become perceptually salient is importantly related to questions of musical (dis)similarity judgments. It has been proposed that the two are, in fact, symbiotically related: (dis)similarity judgments are derived from comparing perceptually salient entities while at the same time, entities become perceptually salient via (dis)similarity judgments.

#### 4.2.1.1.1. *Cues and structure*

Irene Deliege's seminal work addresses how entities gain perceptual salience through (dis)similarity judgments, thus shaping our sense of a work and its structure. Deliege's research suggests that "group boundaries appear spontaneously when a contrasting element stands out from the more undifferentiated environment" (2007, 15). Thus, the perception of both similarity *and* difference enable and are required by our sense of grouping and segmentation within an auditory stream. Such ideas are the basis for a model based on the principles of similarity and difference: "*similarity* proposes that small differences between elements within a unit are minimized, and *difference* assumes that the contrast between elements adjacent to a boundary is emphasized" (Deliege, 2001, 236; also see Deliege, 1989). This observation inspired a model where an attentive listener comprehends *cues*, which are "brief but meaningful and significant structures, which stand out from the sound background" (13). These cues help listeners create a representation of the piece's structure. The perception of structure and form is therefore governed by the aforementioned principles of similarity and difference: "the principle of sameness ... constitutes groups and groupings of groups; the principle of *difference* ... differentiates them" (Deliege and El Ahnmadi, 1990, p. 20). Cues are specific to each piece, they are often picked up subconsciously and importantly to this discourse, they

stand out on the basis of (dis)similarity.

To explore these ideas, Deliege and El Ahnmadi (1990) conducted an experiment involving perceptual segmentation boundaries that we can use as a model to explore musical (dis)similarity and (dis)continuity. Specifically, the authors asked participants to indicate segmentations in Berio's *Sequenza VI*. The authors predicted that participants would indicate segmentations that correspond to moments of high contrast. The concluded that:

Any rest followed by the continuation of the same cue could be reinterpreted as an internal break within the section. On the contrary, the absence of the cue alerts one's attention toward a change, a turning point in the formal organisation. The more the contrast is clear cut, the more the break is perceived as important, whatever the real length of the rest might be. (42)

Thus, perceptual segmentation boundaries correspond to moments of change, or, moments of dissimilarity and discontinuity. The higher degree of change, or, the more dissimilar / discontinuous the moment, the more salient the perceptual segmentation boundary will be. We can implement this methodology in a study of discontinuity by first asking participants to identify perceptual segmentation boundaries and then analyzing the entities adjacent to such boundaries to discover what qualities are most salient and relevant in that particular musical context.<sup>31</sup> We can then compare the relative (dis)similarity of these qualities in order to describe the type and degree of discontinuity present.

In regard to the question of conscious willing in these matters of attention, Deliege distinguishes between *implicit* and *explicit* perceptual processes. The former refers to the listener not being aware of having access to the phenomenal information. The latter refers to the listener willfully seeking phenomenal information. In order to examine the extent to which similarity relations play an *implicit* role in musical memory,

---

<sup>31</sup> See Deliege (2007) for more detail in regard to how dis(similarity) perception and the comprehension of cues affects our structural understanding of a piece.

experiments were conducted that asked participants to 1) identify structural segmentations in musical works and 2) after hearing a piece of music, locating segments of that music in the temporal order originally heard. The experiments suggest that during the listening process, cues are oriented implicitly within a mental schema. Further, the representations that result from these cues are more strongly linked with resemblance versus reasoned, abstract categories. We can thus distinguish between *perceptual cues*, based on appearance, and *conceptual cues*, which are related to meaning (Deliege, 2007, 19). In order to examine the extent to which *explicit* processes play a role in similarity perception and musical memory, experiments were conducted where participants had access to the data acquired during the preparatory phase, and the instructions to them identify similarity / resemblance as the key reference for their answers. Participants were asked to rate the frequency of reference motives, classify motives into groups, assess the degree of similarity of pairs of motives, and memorize a motive and later recognize it in a musical section. Results showed that explicit recognition processes were reliable for both musicians and nonmusicians (Deliege, 2007, 21-27). Deliege concludes that while the impact of similarity relations was greater when instructions were explicit, “all results showed the impact of similarity” (2007, 29).

Deliege’s work illuminates important aspects of our perceptual experience of music and musical (dis)similarity, thus, some of her theoretical and methodological approaches should be incorporated into the current framework. The idea of identifying cues and/or structural boundaries in order to illuminate the most salient points of difference is useful in experimental, analytical and compositional methodologies. Identifying and rating structural boundaries and cues can be done without deconstructive analysis, and thus presents a potentially fruitful approach to describing similarity relations holistically.<sup>32</sup> Both can also give us a sense of the qualities of and

---

<sup>32</sup> One could argue that there are other factors outside dissimilarity relations that create structural boundaries in music. One must be careful therefore to identify such factors when discussing (dis)similarity.

relationships between entities within a piece. When relational difference promotes a perceptual cue or boundary, we gain insight into the qualities an entity must exhibit to be outstanding, which subsequently gives us a sense of the qualities of the surrounding context. This latter point, and the notion of context, is particularly important to (dis)similarity judgments.

#### **4.2.1.2. Categorization**

Categorization is another cognitive process that is derived from and influences (dis)similarity percepts. Deliege (2001) and Cambouropoulos (2009) describe how categorization processes play an important role in (dis)similarity judgments and musical experience in general. Indeed, the notion that similarity and categorization are connected seems reasonable: “The idea that similar entities tend to be grouped together into categories is intuitive, and is supported by data collected by Rosch and colleagues in the 70s (*e.g.* Rosch 1975, 1978) (Cambouropoulos, 2009, 10). If we accept this connection, how then do our mental categorizes affect our (dis)similarity percepts?

Cambouropoulos considers the interactions of these processes:

It is not simply the case that one starts with an accurate description of entities and properties, then finds pairwise similarities between them and, finally, groups the most similar ones together into categories. It seems more plausible that as humans organize their knowledge of the world, they alter their representations of entities concurrently with emerging categorizations and similarity judgments. Different contexts may render different properties of objects/events more diagnostic concurrently with giving rise to certain similarity relationships and categorisations. If context changes, it affects similarity, categorisation and the way the objects/events themselves are perceived (Cambouropoulos, 2009, 11).

Categorization, then, is not just the by-product of similarity judgments. Instead, our mental categories shape both the entities that we perceive and the comparisons that we make between those entities. The framework developed in this project recognizes the importance of categorization processes in that (dis)similarity judgments both reflect



current categories and shape categories that result from musical experience.

Categorization can also help explain tendencies that are exhibited in experimental data. For example, perhaps a particular respondent draws structural boundaries where guitar timbre changes. If that person is an experienced guitarist that has developed rich and distinct categories of guitar timbres, we could better understand how those individual aptitudes and experiences inform (dis)similarity judgments. From there, we can see to what extent these judgments are shared (or not) by other participants, giving us some insight into the commonalities and differences between different people's (dis)similarity judgments.

As analysts and composers, we must therefore be sensitive to categorization's role in musical (dis)similarity and (dis)continuity. As musical analysts, we must consider to what extent an entity represents a larger category that may carry with it a diversity of musical and non-musical associations. As composers, we can make similar distinctions, placing certain entities in a musical texture as a function of their categorical associations. In both cases, we can recognize that entities can evoke categorical representations to varying degrees, that is, an entity could be a categorical exemplar or could merely give a hint of a mental category. The degree of categorical association thus is a parameter that we can use in our analytical and compositional efforts.

#### **4.2.2. Musical Context**

As previously stated, one of the main assertions of this paper is that musical context influences (dis)similarity judgments. Emiliós Cambouropoulos' contemplations of this idea led him to the conclusion that "similarity always depends on context (*i.e.* it is contextually defined)" (2009, 8). He illustrates this by considering Ahlback's "50% threshold idea", which essentially states that two entities are similar if they have more than 50% of attributes in common and dissimilar if they have less than 50% of attributes in common (2009, 16-18). Cambouropoulos points out that such an objective measure is

problematic because it doesn't consider context: two entities could share few features but still be considered similar if the context was particularly diverse, while two entities could share many features but still be considered dissimilar if the context was homogenous (2009, 17-18). Cambouropoulos (2001) illustrates this idea via the following musical examples (Figure 4.8):



**Figure 4.8** Excerpts from Steve Reich's *Electric Counterpoint* (top) and Iannis Xenakis' *Keren* (bottom) (Cambouropoulos, 2001, 357).

He hypothesizes that if we were to ask subjects to choose the pair that is most similar, some would choose the top pair and some would choose the bottom. He points out that in context, the relative similarity of the excerpts is perceived differently: the top pair is perceived as dissimilar in the "homogenous" context of Steve Reich's *Electric Counterpoint* while the bottom two are perceived as similar within the "diverse" context of Iannis Xenakis' *Keren* (357). Thus, we cannot say that some level of commonality between two entities will necessarily result in notion similarity, for the latter depends on level of commonality *given a particular context*. Instead, we must consider what level of commonality is necessary to produce notions of (dis)similarity in a particular context, which we can then use to guide our evaluation of commonality between entities and relations.

By identifying the relevant components and (dis)similarity requirements of a

particular context, we can start to address some of the issues that have been identified as problematic for similarity as a construct, such as accounting for attentional, cultural and individual differences. Asking participants, or ourselves as analysts and composers, about salient musical cues and/or structural boundaries seems an effective way of identifying the relevant features and the degree of commonality required for similarity within a specific musical context. We then can start to explore why, in terms of music and psychology, people exhibit differences and commonalities in their (dis)similarity percepts. We want to know about both: we recognize that our musical percepts can differ (for any or all of the aforementioned reasons), and we want to know the reasons for those differences. At the same time, we have a strong intuitive sense that our individual experience of a piece of music corresponds to that of other people (the exploration and elucidation of this sense is an important part of group listening, both inside and outside educational contexts). This sort of experimental methodology could help reveal these perceptual commonalities. As described in the previous paragraph, such cues / boundaries can inform our musical experience whether they are obtained explicitly or implicitly, so we must be mindful that we can ask such questions indirectly or directly in our explorations.

## Summary

Our (dis)similarity percepts influence our experience of musical (dis)continuity, and music in general. They allow us to segment a musical surface, extract the most salient entities and make associations between those entities. (Dis)similarity judgments involve the perception of both likeness and difference; the proportions and interactions between these likenesses and differences determine the degree of (dis)similarity between the compared entities. A number of conceptual approaches have been developed to describe these proportions and interactions including spatial, featural, structure-

mapping, and transformational models. This project proposes a new approach that incorporates these various models when comparing the (dis)similarity of entities.

When we describe musical (dis)similarity, we must consider not only the qualities found *within* compared entities, but also the musical and psychological context in which compared entities occur. In regard to psychological context, (dis)similarity percepts both guide our attention and result from its focus. Irene Deliege's seminal work addresses how entities gain perceptual salience through (dis)similarity judgments, thus shaping our sense of a work and its structure. We can use Deliege and El Ahnmadi's (1990) experimental methodology involving perceptual segmentation boundaries as a model to explore musical (dis)similarity and (dis)continuity. Categorization is another cognitive process that is derived from and influences (dis)similarity percepts, and thus should be considered when describing or creating musical (dis)similarity and (dis)continuity. Musical context, the textures and ideas that surround the compared entities, also significantly influence the perception of relative (dis)similarity, so we must be careful to judge the (dis)similarity of entities in their appropriate contexts. By identifying the relevant components and (dis)similarity requirements of a particular context, we can start to address some of the issues that have been identified as problematic for similarity as a construct, such as accounting for attentional, cultural and individual differences.

As mentioned in Chapter I, because musical entities, attributes and relations vary widely, it is not unreasonable to describe them with a variety of tools, some of which have been mentioned in this chapter. An explanatory / predictive framework must be flexible enough to deal with such diversity if it is to say something meaningful about musical perceptual experience. One may object that it is impossible to make apples to apples comparisons when using different tools or methodologies. The response is that a variety of tools and methods can co-exist in a framework, but they need to be specified and weighted appropriately depending on the particular (dis)similarity comparison.

Such a system of weighting would allow one to derive representations of wholes that can be used in relational comparisons.

## 5. Chapter V: Parts, Wholes and Hierarchies

As we bring entities into relations, we can perceive new kinds of organizations. In chapter III, I described how we relate individual pitches to form intervals, which we then relate to form a continuous ascending gesture. In that same chapter I discussed how we connect temporally-separated entities on some level beyond the musical surface to create *interrupted continuity* (on the musical surface, the constituents in an interrupted continuity relation are not connected by definition). Indeed, our ability to abstract and compare entities to create new perceptual organizations is a foundational part of how we experience musical (dis)continuity. We must therefore consider the components and implications of this process in more detail.

We can describe entities, relations and the organizations they inspire in the language of parts, wholes and hierarchies. As previously described, parts, such as individual notes, can form a whole, such as a rhythm. We can specify the identity of these parts by looking at them either inside or outside the whole that they constitute (I will explore this topic in the first part of this chapter). As we perceptually combine parts to form a whole, the latter can project emergent properties.<sup>33</sup> We can then specify the identity of wholes via their emergent properties, such as the *tempo* of a rhythm (emergent properties lend to parametric description, which will be the subject of the next chapter). We can illuminate the relationship between wholes and parts by describing them relative to a hierarchy where larger and more abstract organizations inhabit higher hierarchical levels. The hierarchical levels that entities and relations inhabit affect our (dis)similarity and (dis)continuity percepts in various ways. The possibilities of such interactions will be the focus of the bulk of this chapter.

---


<sup>33</sup> Rock (1986) defines an emergent property as “one that arises only by the manner in which parts relate to one another and that therefore cannot possibly inhere *in* the parts” (33-3).

## 5.1. The identity of parts

We can describe parts in isolation or from the perspective of the whole(s) that they constitute. In terms of the former, we can think of parts as being causally prior to and isolatable from wholes. We can then identify the specific qualities of those parts, which we can incorporate in our descriptions of (dis)similarity and (dis)continuity. In contrast, Rock (1986) discusses the Gestaltist position that we understand the identities of parts through our experience of the wholes that they constitute. In order to reconcile these seemingly contrasting positions, Rock suggests that we distinguish between two meanings of “part”:


First, the properties of a part or unit qua proximal stimulus must be detected; for example, the slope of a line’s image ... The whole is then constructed on the basis of the interrelationship of all such proximal stimulus parts ... Second, certain phenomenal properties of parts arise as a function of the emergent properties of the whole configuration. (1986, 33-11)

Thus, parts can project different properties depending on whether they are viewed in isolation or in the context of a whole that they constitute. We can distinguish between these two meanings of “part” with the terms *proximal stimulus properties* and *phenomenal properties* when discussing a particular organization. We must be careful to not confuse these meanings when we consider parts in isolation (we must do the latter if we want to give detailed descriptions of music) for the identity or function of a property (or entity such as a melody, rhythm or timbre) can differ depending on context. For example, consider two sounds separated by a temporal interval, 200 msec in the case of Figure 5.1.




**Figure 5.1. A 200 msec sonic interval.**

Such an entity gains a different identity when experienced in the context of a larger emergent rhythmic organization, such as in Figure 5.2:



**Figure 5.2. Rhythmic Configuration with 200 msec interval.**

The same interval projects yet another identity if we situate it in a different context, such as in Figure 5.3:



**Figure 5.3. Another example of a rhythmic configuration that includes the 200 msec interval.**

In the language of Rock, we can understand Figure 5.1 as a *proximal stimulus property*, given it, in relation to other proximal stimulus parts, constitutes the larger whole. Alternatively, Figure 5.2 and Figure 5.3 illustrate *phenomenal properties*, which are a function of the emergent qualities of the rhythmic configuration as a whole. Thus, the same part can be interpreted in a number of different ways depending on its context. Given that context, as discussed in the previous chapter, plays an important role in the way that we perceive musical (dis)continuity and (dis)similarity, our diligence in distinguishing between these two properties will help us understand (dis)continuity.



## 5.2. Perceptual hierarchies

### 5.2.1. Overview

We can think of parts and wholes as functioning within a hierarchy, where parts are lower-level and wholes are higher-level. The levels within a musical hierarchy include (from lowest to highest) individual constituents (such as single notes), groupings and relationships formed by individual constituents (such as an *interval* between two notes), groupings of these groupings (themes, melodies, rhythms, phrases), and so on to larger organizations that include sections, movements and works. Our descriptions of entities' attributes can be oriented relative to this hierarchy. For example, we can describe attributes that correspond to lower-level entities, such as the pitch or duration of an individual note or we can describe attributes that correspond to higher-level entities, such as the tempo of a rhythm.<sup>34</sup> Relations can similarly be described hierarchically. There can be a relation between two adjacent lower-level entities, such as an interval formed from two adjacent pitches, a relation between medium-level attributes, such as between the pitch sequence and rhythm that constitute a phrase, or a relation between higher-level entities, such as between a bass line and a vocal part. The sense of abstraction increases as we move up the hierarchy, that is, the more encompassing the group, the further we move away from the individual sonic events. By identifying a whole's attributes and relations on various hierarchical levels, we form a basis from which we can compare the whole to other entities that have (or don't have) corresponding attributes and relations, as described in chapter IV.

### 5.2.2. Hierarchies and Attention

We are able to attend to a number of different hierarchical levels during a

---

<sup>34</sup> Many of our musical concepts are descriptions of the emergent identities of these various groupings, such as intervals, rhythms and melodies.

listening experience; the level that we attend to influences our (dis)continuity percepts. For example, consider a 3:2 son clave-based rhythm, such as a slow mambo (Figure 5.4):



**Figure 5.4. Slow Mambo. Brown. *Afro-Latin Rhythm Dictionary: Handy Guide* p.13.**

immediately followed by a waltz, such as the one found in Strauss' *Voices of Spring* (Figure 5.5):



**Figure 5.5. Johann Strauss' *Voices of Spring*. Arranged by Sebastian Spychaj.**

The lowest hierarchical level includes organizations that result from the relations between individual notes or sounds. As a listener attends to this level in the context of rhythm during *Voices of Spring*, she perceives continuity starting with the third beat of the first measure through the first beat of the second measure (indicated by the purple bracket in Figure 5.5), at which point an interrupting rest effects a sense of discontinuity. If the listener attends to groupings of individual elements into larger organizations such

as phrases, bars and meters, she could perceive (dis)continuity on higher levels. For example, if the two figures above were juxtaposed as described, a higher-level rhythmic discontinuity would be created as a function of conflicting time signatures, accent patterns and temporal configurations. As a listener attends to even larger groupings (such as sections and movements), (dis)continuity can be perceived on even higher levels.

### 5.2.3. (Dis)continuity on Multiple Hierarchical Levels

Different hierarchical levels can describe how different kinds of (dis)continuity can occur simultaneously. For example, imagine a person who washes the dishes ('A1'), stops to read a book ('B'), and then washes more dishes ('A2'). The higher-level activity 'washing the dishes' could be considered as an example of interrupted continuity: the process involved in 'A1' is sufficiently similar to the process involved in 'A2' and thus 'A1' and 'A2' are perceptually connected and thus, continuous. At the same time, the higher-level category 'washing the dishes' contains lower-level activities such as 'washing a particular plate'.<sup>35</sup> If a particular plate was fully washed and identified as activity 'A1', we would not say that activity continued into 'A2', where different plates were washed. From this medium-level perspective, activity 'A1' (washing a particular plate) is interrupted by discontinuity 'B1' (reading a book), which is subsequently interrupted by discontinuity 'C1' (washing a different plate).<sup>36</sup> As a result, continuity is perceived on the highest-level while discontinuity is perceived on lower on a lower level.

To illustrate these ideas in music, consider the following example (Figure 5.6):

---

<sup>35</sup> Note that 'washing a particular plate' could be considered a continuous entity itself, showing how different types of continuity can also functionally coexist in hierarchical organizations.

<sup>36</sup> This representation seems intuitively strange because we tend to group the activity 'washing a plate' together with other like activities as according to higher-level representations such as "washing the dishes".



**Figure 5.6. Interrupted Continuity.**

In one sense, consistent with the type of interrupted continuity already described, we can consider each three-note group as a continuous group and each rest a separate entity that promotes discontinuity. Thus, we can represent the above as 'A1' 'B1' 'A2' 'B2' etc., where the 'A's are connected despite the interruption from the 'B's. Alternatively, we can combine each three-note group with each subsequent rest to create a higher-level group whose duration is a quarter note. The example would then be represented as 'A1 A2 A3 A4'. Perceptually, this means that we experience four adjacent identical entities. The sameness of each entity creates a connection between them, thus inspiring a sense of continuity. This continuity is different from the lower-level kind that experiences distinct three note groupings separated by rests. On this higher-level, our sense of continuity is temporally expanded (it encompasses all four 'A' blocks) and it is more abstract in that it 1) is further distanced from the sonic phenomenon and 2) both continuity and discontinuity exist within its lower-level constituents. It is thus 'higher-level' and different in kind as compared to the aforementioned 'interrupted continuity' represented by as A1 B1 A2 B2.

To further illustrate how (dis)continuity can function on multiple hierarchical levels, consider Figure 5.7 and Figure 5.8:



**Figure 5.7. Low-level discontinuity, higher level continuity.**

In the first measure of Figure 5.7, the lower level continuity created between each three

note group (as in Figure 5.6) is subsequently disrupted by the triplet eighth groups in the second measure, creating lower level discontinuity. At the same time, the higher-level continuity created in the first measure is *maintained* in the second measure as this continuity requires the abstract configuration *entity + rest; duration - quarter note*, which both measures fulfill.

In Figure 5.8, Figure 5.6, stated in the first measure, is subsequently followed by the rhythm of Figure 5.5 (*Voices of Spring*).



**Figure 5.8. High and low-level discontinuity.**

As in Figure 5.7, the rhythm of the second measure disrupts the continuity created by the grouping of the three sixteenth notes, creating a lower-level discontinuity. At the same time, we also have a sense that some other kind of sequence / gesture / grouping has been interrupted. This different sense of discontinuity results from the disruption of a different kind of continuity (a higher-level one). In this case, the latter is characterized by the configuration *entity + rest; duration - quarter note*, which is established in the first measure. It follows that continuity on each hierarchical level has a corresponding type of discontinuity.

### **5.2.3.1. Sequences, Processes**

Thus, when we group discrete lower-level components, we are able to perceive higher-level continuity that can be described in terms of ideas such as ‘sequences’ and ‘processes’. Such insight allows us to better penetrate the Oxford English Dictionary’s second entry for the definition of ‘continuity’:

Of immaterial things, actions, processes, etc.: The state or quality of being uninterrupted in sequence or succession, or in essence or idea; connectedness, coherence, unbrokenness (oed.com).

This definition highlights how continuity can be perceived among higher-level organizations such as ‘sequences’, ‘successions’, ‘essences’ or ‘ideas’. This is a more abstract and more flexible version of the idea of continuity than the one first defined in Chapter III as “the state or quality of being uninterrupted in extent or substance” (oed.com). Here, the salience of lower-level disjunctions fades as a perceiver apprehends higher-level organizations. One is able to bring these higher-level entities into relations and is able to conceive of connections, processes and continuities to describe them.

These sequences / processes typically involve an entity and some operation that is applied to that entity to produce a subsequent iteration: I will refer to *entity + operation* as a *transformational equation*.<sup>37</sup> To illustrate this idea, consider Figure 5.9:



**Figure 5.9. Hierarchical grouping.**

The *entity* part of the equation is defined by how we group components. Here, as in Figure 5.6 and Figure 5.7, one could group individual adjacent elements together so that three-note groups are interrupted by eighth note rests, as depicted by the orange brackets in Figure 5.9. Each of these three-note groups exhibits some degree of continuity, given the temporal and pitch relationships (adjacent and one scale step apart,

<sup>37</sup> Sequences and processes are different but related ideas. A process in this context refers to the steps that must be actioned to achieve a particular end. A transformational equation describes those steps. A sequence is a set of entities that are related in a particular way. A transformational equation describes how those entities are related. Thus, we can think of processes as generating sequences; transformational equations describe them both.

respectively) of the constituent elements. From this perspective, the passage would be perceived as an example of interrupted continuity. Alternatively, if a listener were to group each three-note sequence and the subsequent rest together (depicted with blue brackets in Figure 5.9), we could represent the passage as four groups lasting a half note each, as in the case of Figure 5.6.

The *operation* part of the equation involves how the entities relate to each other. Unlike Figure 5.6, the pitch content of each group in Figure 5.9 differs from its neighbors, so one may be tempted to represent the sequence as ‘A B C D’. The latter is problematic though, given that a listener would likely identify a relationship between each three-note group, specifically that each group ascends the C Major scale by step, and that the first note of each group is one scale degree higher than the first note of the previous group. In this way, each group is intrinsically connected to its predecessor: each group is its *predecessor (entity) + operation* (the operation in Figure 5.9 is that the group is shifted up one scale degree). This abstract transformational equation articulates the *relation* between entities in the sequence. Because each adjacent group shares the same relation, we perceive similarity between the groups. This similarity inspires a notion that these entities share a common higher-level *identity*<sup>38</sup> and are thus strongly connected, which gives inspires a notion of continuity. We typically describe these types of similarity / continuity percepts in terms such as ‘sequence’ and ‘process’.

The mental abstraction and relating just described also explain how we group and perceive pitch continuity on the lowest level (orange brackets). Within each three note group, the subsequent pitch is derived from the same transformational equation *predecessor + operation* (in this case, the operation is to add one scale degree to the previous pitch). As in the case of higher-level continuity described above, the equation articulates the relation between entities. As adjacent entities share the same relation, we

---

<sup>38</sup> Diana Deutsche notes that in 1890, Von Ehrenfels described our ability to perceive auditory sequences as retaining their identities despite transposition (as long as inter-sequence relationships are maintained) (32-25). Here, the example involves shifting a sequence within a scale rather than transposition, but the basic idea is the same.

perceive similarity between groups that then inspires a notion of continuity. This example also shows the abstractness of transformational equations; here, the same equation creates two different kinds of continuity on two different levels.

Because sequences and processes are derived from transformational equations, we can create discontinuity by introducing a break / gap / interruption in the course of such sequences and / or processes. As stated earlier, we can do this by creating dissimilarity between operations and/or dissimilarity between operands. In the example of Figure 5.9, the former could be achieved by transposing a number of subsequent groups up one scale degree, as occurs in Figure 5.9, and then transposing a group down two scale degrees, breaking the sequence. In regard to the latter, we could append the group G E C# to the end of Figure 5.9, which would break the adjacent scale-step configuration of the previous groups, and thus would interrupt the sequence.

There is a connection between the preceding discussion and the one featured in chapter III involving ascending pitch sequences; in both cases, I describe how we draw abstract relations between distinct entities from which we form (dis)similarity and (dis)continuity judgments. In chapter III, I described the resulting percept in terms of a general “ascension”. Here, I described the resulting percept in terms of a transformational equation. One may then ask, do we perceive specific transformational equations or a more general sense of movement? The answer depends on listener knowledge / aural skills (context) and level of hierarchy.

#### **5.2.4. Hierarchies and Context**

Context, which comprises individual differences, affects how an entity will be perceived relative to a hierarchy. If a listener is unfamiliar with a major scale, then she wouldn't be able to perceive the specific operation (adding one scale degree) involved in a transformational equation. With that said, assuming that such a listener can make



registral distinctions, she *can* understand the relation between adjacent pitches (each subsequent pitch is higher than the previous one). This relation is derived, consciously or unconsciously, from the understanding of an abstract kind of transformational equation defined by *predecessor + ascending transposition*. Given this equation produces the various constituent intervallic relations, the relations share an inherent sameness. This sameness, as described in chapter III, inspires a notion of continuity, which in this case, is defined by ascension.


If a different listener *is* familiar with and *can* identify the major scale (*structural knowledge*), then she *could* perceive the operation and thus the equation. She then perceives relations not only in terms of “ascending” (which she still does perceive), but also in terms of the precise equation that defines this particular type of “ascension” (predecessor + a scale degree). As this understanding is a function of a larger complex of structural knowledge that involves intervals, interval sequences, harmonic function, etc., we can consider it higher-level than a more nebulous sense that doesn’t involve this rich network of information. Given these different percepts, which differ because of context, correspond to different hierarchical levels, they can produce different notions of (dis)continuity. Context thus affects hierarchical perception.<sup>39</sup>

We can apply the above discussion of how context and hierarchies affect the perception of continuity to other musical concepts including (but not limited to) rhythm. For example, in the case of the first measure of a 3-2 son clave rhythm (measure one of Figure 5.10), one listener may have a general sense that a particular pattern is syncopated while a more expert listener has both a general sense of syncopation and can locate specifically where the syncopation occurs (the second note, which is located between beats two and three of the first measure). The latter percept, as in the earlier example of the major scale, is enabled by a rich complex of structural knowledge that

---

<sup>39</sup> There is an intuitive attraction to this idea given as musicians, our sense of directional pitch sequences varies depending on whether or not we are familiar with the underlying structure governing such movements. The two ascending pitch sequence examples in chapter II perhaps highlight this distinction for some listeners.

involves beats, beat division, tempo and meter.



**Figure 5.10. Two examples of syncopated rhythms.**

If we reconfigure the pattern so that the second note occurs on a beat while the third note falls between the third and fourth beat, we can effectively shift the syncopation (second measure of Figure 5.10). Were we to juxtapose these two examples, an expert listener could potentially experience such configural change as a break in continuity, given his structural knowledge. Conversely, a listener who only has a general sense of syncopation may perceive continuity between both these configurations in that they are both syncopated. Importantly, the less expert listener can perceive difference between the patterns (without being able to identify the specifics of said difference) and the expert listener can experience the more general continuity associated with syncopation (despite her knowledge of the specific differences between the rhythmic configurations). One explanation of why these various precepts can simultaneously occur is because they function on different hierarchical levels (the more specific rhythmic interpretation can be thought of higher-level given it involves a rich network of structural knowledge, as in the case with pitch as described earlier).<sup>40</sup> The ability of a listener to perceive (dis)continuity on various hierarchical levels is partly a function of individual differences (context).

Thus, we must be sensitive to how context and hierarchies affect the perception of (dis)continuity. Different listeners distinguish different features within a musical texture because of contextual factors. Each of these distinctions evokes a set of

<sup>40</sup> Experimental evidence of such simultaneously occurring but contradictory percepts would strongly support the idea that we experience music hierarchically, for how could contradictory percepts be simultaneously salient if the objects of those percepts were inhabiting the same level?

associations that depend on musical training, personal preferences, stylistic familiarity, etc. These various distinctions and associations gain perceptual salience on different hierarchical levels, which produce various kinds of (dis)continuity. Thus, as listeners differ in their percepts because of contextual factors, they will perceive different kinds of (dis)continuity within a piece. Experimentally affirming that these different scenarios produce different kinds of continuity judgments would strengthen the requirement that we view (dis)continuity flexibly by remaining mindful of the influence of contextual factors.

### 5.2.5. Hierarchies and Similarity

Context that inspires higher-level configurations can affect the perceived (dis)similarity of entities. For example, most would judge the sequence of sounds in Figure 3.12 to be highly dissimilar.

[Click to listen](#)

**Figure 5.11. The segmented sounds of Girl Talk's *Pump It Up*.**

In the context of Girl Talk's *Pump It Up* though, these disparate sounds combine to form a higher-level emergent rhythmic organization (Figure 3.13).

[Click to listen](#)

**Figure 5.12. Excerpt from Girl Talk's *Pump It Up*.**

Most would agree that Figure 3.13 is significantly more continuous than Figure 3.12. This continuity, as described previously, is higher-level as it is created from relational similarity between lower level constituents: lower-level constituents relate to the larger emergent whole in similar ways. It also may be the case that because of this higher-level continuity, we judge the qualities of constituent entities as more similar than we would

in the more temporally separate configuration as in Figure 3.12. It is also possible that our judgments of the similarity between distinct entities are influenced by the hierarchical level that those entities inhabit. For example, the way that we judge the similarity between movements may be different than the way that we judge the similarity between timbres of individual notes. If this were the case, we would need to adjust our explanatory models accordingly, or develop new ones. While these possibilities need to be experimentally tested, they at least caution us to avoid evaluating similarity out of hierarchical context.

#### 5.2.6. Hierarchies and Transformations

We also must consider the hierarchical level that is most perceptually salient when describing transformations. For example, imagine we are listening to a solo flute that, at some later point, is processed with a chorus effect. Perhaps this particular chorus algorithm is complex, a fact that could be illuminated if we were to approach the transformation through an audio analysis methodology. The result could be a fairly extensive list of instructions that is required to transform A to B. For one familiar with the chorus effect though, one could describe the perceptual experience a single instruction: “add chorus effect to the flute”. Indeed, the latter seems to better capture our perceptual experience; it is also likely a better description of the compositional process as well. The chorus effect is thus a higher-level entity that comprises individual qualities, such as *speed* and *depth*, which are derived from even lower-level attributes such as the algorithms that specify delays and modulations. Just because these lower-level attributes are physically present does not mean that they are psychologically meaningful. If we are to discuss transformations, and we should given their descriptive power, then we must make sure that our level of description is aligned with the hierarchical quality of an entity that is perceptually relevant.

### 5.2.7. Varied Saliency Among Hierarchical Levels

If we can perceive entities on multiple hierarchical levels, one wonders which perceptual level we favor, or which we perceive first. In the case of continuity between distinct entities, do we notice fine qualitative detail in both entities and then compare them on a point-by-point basis to arrive at a general notion of (dis)continuity or are we struck by some notion of (dis)continuity on a higher level, from which we can look at the entities in more detail and then figure out what aspects caused the judgment? Are certain levels of a hierarchy more perceptually salient than others?

There is support for the notion that higher-level musical organizations are more perceptually salient than their lower level constituents. Music researcher and theorist Emilios Cambouropoulos (2009) offers the following:

There is evidence that things such as melodic and harmonic pitch intervals, chords or larger configurations such as tone clusters, tremolos, trills, glissandi are commonly perceived by listeners as wholes rather than combinations of atomic lower-level components. For example, especially for pitch, it has been suggested that the majority of listeners, for whom musical pitch is relative, perceive pitch intervals categorically prior to individual pitches (Dowling and Harwood, 1986; Handel, 1989). Tenney suggests that larger sound complexes such as tone-clusters or other dense chords “cannot usually be analysed by the ear into constituent tones, and [he suggests] are not intended to be analysed.” (Tenney, 1961:6) — see also (Cook, 1990); even simpler triadic chords may be perceived as elementary chord types — or even tonal chord function types — before being analysed into their constituent tones and intervals. A glissando is also perceived and can be represented as a single entity (see Bregman 1990, p. 644) with start-pitch and end-pitch, slope of transition, duration and intensity (a linear transition between the two pitches may be implied as a default) (p. 13).

In an experiment involving similarity judgments of Western Popular music, Novello et al. (2011) found that genre, tempo and timbre showed “a hierarchical degree of impact on participants’ pair rankings (genre > tempo > timbre)” (1). The authors point to other studies that support these findings including Chupchick et al. (1982), Lamont & Dikken (201) and McAdams & Matzin (2001). Thus it seems we sometimes privilege higher-

levels over lower-level ones when listening to music.<sup>41</sup> If this is indeed the case, then we need to weight higher-level organizations appropriately if we are to describe the perception of (dis)continuity accurately.

These observations are also in agreement with earlier discussions in this paper, suggesting that the identity of configurations is defined more by higher-level qualities, such as the relationships between the constituents, then by the lower-level qualities of the constituents themselves.<sup>42</sup> Consider an ascending pitch gesture: the defining characteristic of this organization is a function of the relationships between constituents, specifically, that each pitch is registrally higher than its predecessors. This one simple rule ensures that all components have a specific relationship with near and distant pitches, which governs the overall shape of the configuration. The actual pitch of each component plays a less important role in producing this organization; we could change any or even all of the pitches in the sequence and, as long as the proper relationships are maintained, create an ascending gesture. Further, as shown in chapter II, the interval quality between each adjacent pitch also plays a less important role in creating this organization: as long as the general relationship between each group of adjacent pitches (specifically, where the latter is registrally higher than the former) an ascending gesture will be perceived.<sup>43</sup>

Our musical experience further supports the notion that we privilege higher-level discontinuities over their lower-level counterparts. When a jazz section is juxtaposed with a noise-rock section, as is the case in Naked City's *Snagglepuss*, the difference in genres is more striking than the difference in hi-hat patterns. In Coffey's *Koans*, our first reactions are "recordings of street noise are juxtaposed with electronic noise elements" and not "the spectral character of *entity A* contrasts with that of *entity B*" or "*entity A*

---

<sup>41</sup> Such preferences undoubtedly depend on the character of the individual elements, higher-level groupings and overall musical context. Further research could examine the strength of these tendencies.

<sup>42</sup> Gentner (1983) supports the idea that relations are more salient than object-attributes in analogy.

<sup>43</sup> Of course, the qualities of the constituent pitches and intervals will influence the character of the ascending gesture. The point here is that we can alter the pitches and intervals and still maintain the general sense of ascension.

contains little pitch information while *entity B* contains the pitch set [...].”

There are several potential explanations for why this is the case. Perhaps we prefer more general and abstract concepts that are singular and simple versus the dense network of information that constitutes such concepts, which is plural and complex. The former requests fewer processing resources. Perhaps we understand details by first understanding generalities: without the more general category, we may not make connections between the individual instances. We can best relate and differentiate between flute timbres when we have a well-established concept of a flute. We privilege high-level categories because they help us understand the details of the world. Perhaps higher-level discontinuities are more salient than lower level ones because they typically contain changes in multiple lower-level parameters. Salience in this case is simply a matter of the amount of change, rather than any special qualities endowed by particular hierarchical levels. Perhaps high-level discontinuities, as articulated by *Snagglepuss*, are more noticeable because they are more rare than drastic juxtapositions along lower-level musical parameters. Such uniqueness demands our attention. This would support the notion that the perception of high-level organizations and/or discontinuities depends on listener knowledge. We therefore cannot assume that two different listeners will necessarily gravitate towards the same organization and/or discontinuity because it is high-level, for the latter is contextually defined. The reality of the situation likely involves some combination of these positions.

### 5.2.8. Non-Hierarchies

Describing music in terms of hierarchies is useful in that it provides a way to illuminate underlying structure, which is undoubtedly why many analytical writings and methodologies have embraced such representations. However, we should be careful not to assume such representations express the totality of our musical experience. In many ways, the experience of listening to music isn't hierarchical at all. Instead,

notes, melodies, phrases, rhythms, and sections come to us in a way that is *all at once*, wrapped in a particular moment. We can locate all the components of a musical texture on common ground, where various kinds of groupings interact in an immediate and available way. When a performer conveys deep joy in his playing, that joy is inseparably intertwined with the spectral and dynamic qualities of the sounds that evoke such associations. We can separate and assign the various components of such expression to a segmented hierarchy, but such hierarchical representation does not illuminate the magic of this common ground. Thus, if we are to use hierarchical descriptions, which seem fully justified, we must keep our mind open to the possibility that there are some perceptually salient aspects of musical experience that may require other descriptive means. Indeed, we may be clued in to this necessity if our hierarchical descriptions of (dis)continuous relations leave something lacking relative to our perceptual experience.

Further, when we start to consider (dis)continuity created through associational juxtapositions, it is not always clear how such percepts fit into a hierarchy. For example, what is the appropriate hierarchical level for discontinuity created via stylistic (in the sense of an improvising performer's "style") juxtapositions? Style can be associated with temporal articulation, timbre, pitch contour, form or extra-musical associations. One could argue that since style refers to multiple levels, it must be a higher-level concept in the same way that a section refers to individual notes, rhythms and phrases. One could say that since style is a higher-level category formed from a number of individual instances found in a variety of works, the category must be higher-level than a single work. Given higher-level discontinuities are typically more perceptually salient than lower-level discontinuities (the boundary demarcating a section is typically more perceptually salient than the one demarcating a phrase), we could assume that stylistic juxtapositions would be among the most powerful. This could be the case or it could not, depending on the qualities of the juxtaposed styles. If a piano player juxtaposed the



styles of Bud Powell and Bill Evans during a jazz improvisation, the resultant discontinuity may be less salient than a dramatic key change. The salience of stylistic juxtaposition depends on the constituent musical parameters involved and the context. It therefore seems problematic to assign some musical categories to a hierarchy. This doesn't mean that such categories don't function hierarchically, it just means how they function hierarchically may change depending on context.

### Summary

As we bring entities into relations, we can perceive new kinds of organizations. We can describe entities, relations and the organizations they inspire in the language of parts, wholes and hierarchies. We can describe parts in isolation or from the perspective of the whole(s) that they constitute using the terms *proximal stimulus properties* and *phenomenal properties*. We can describe wholes in terms of emergent properties. Together, parts (such as notes) combine to form wholes (such as intervals), which then can combine to form even larger wholes (such as melodies), and so on. These groupings are thus defined by both the number and type of constituents they comprise and the emergent properties that they project. We can represent these parts and wholes as functioning in a hierarchy where particular levels correspond to particular kinds of groupings. In music, groupings can occur on multiple levels simultaneously, and we are able to attend to these various levels in our listening experience. (Dis)continuity can occur on any of these levels, and multiple (dis)continuity vectors can occur simultaneously. These vectors can function antagonistically, so that there is discontinuity on one level and continuity on a higher level. We thus can better understand some of the (dis)continuity types explained earlier in the paper, such as interrupted continuity.

As (dis)continuity can function on higher, abstract levels, it can characterize relational structures such as sequences and processes. We can describe such structures

in terms of *transformational equations*, such as *entity + operation*. As the same transformational equation represents relations between various constituent parts, we perceive those parts to be similar and thus connected in a way that inspires a notion of continuity. Because sequences and processes are derived from transformational equations, we can create discontinuity by introducing a break / gap / interruption in the course of such sequences and / or processes.

Context, which comprises individual differences, affects how an entity will be perceived relative to a hierarchy. If a listener is unfamiliar with a major scale, then she wouldn't be able to perceive the specific operation (adding one scale degree) involved in a transformational equation. If a different listener *is* familiar with and *can* identify the major scale (*structural knowledge*), then she *could* perceive the operation and thus the equation. Thus, different kind of transformational equations are possible depending on context. Different kinds of transformational equations correspond to different kinds of (dis)continuity, which function on different hierarchical levels (entities that are produced from rich complexes of structural knowledge are considered higher-level than their more nebulous counterparts).

Hierarchies also play a role in how we consider (dis)similarity. It may be that higher-level configurations can affect the perceived similarity of constituents. For example, entities that form a higher-level organization may inspire us to view them as more similar than we would if they were perceived separately, as in Girl Talk's *Pump It Up*. Hierarchies also play an important role in transformational descriptions of similarity relations. Given we are interested in the perceptual impact of musical (dis)continuity, we may focus our transformational descriptions on those levels of a hierarchy that are most perceptually salient.

If we can perceive entities on multiple hierarchical levels, one wonders which perceptual level we favor, or which we perceive first. There is support for the notion that higher-level musical organizations are more perceptually salient than their lower

level constituents. These observations are also in agreement with earlier discussions in this paper, suggesting that the identity of configurations is defined more by higher-level qualities, such as the relationships between the constituents, then by the lower-level qualities of the constituents themselves. Our musical experience further supports the notion that we privilege higher-level discontinuities over their lower-level counterparts. If this is indeed the case, we may need to appropriately weight representations of entities that constitute a (dis)continuity relation in order to most accurately reflect perceptual experience.

Describing music in terms of hierarchies is useful in that it provides a way to illuminate underlying structure, which is undoubtedly why many analytical writings and methodologies have embraced such representations. However, we should be careful not to assume such representations express the totality of our musical experience. Further, when we start to consider (dis)continuity created through associational juxtapositions, it is not always clear how such percepts fit into a hierarchy. In cases where hierarchical descriptions do not illuminate the totality of our musical experience, other descriptive means may be required.

## 6. Chapter VI: Parameters and Dimensions

### 6.1. Overview

The preceding chapters, through concepts such as structure mapping, transformational distance, hierarchical representation and contextual influence illuminate how different kinds of relational structures affect our perception of musical (dis)continuity and (dis)similarity. While such concepts are helpful as they illuminate higher-level, abstract characteristics about how we perceive (dis)similarity and (dis)continuity relations, they say little about how different kinds of attributes affect how we perceive structures and the relations they comprise. For example, a certain kind of melodic transformation may be particularly salient and discontinuity-inspiring, while a certain kind of timbral transformation might be less salient and less discontinuity-inspiring. If we are to be able to describe why some kinds of attributes inspire musical discontinuity percepts and some do not, we must be able to describe and represent different types of attributes in detail.

We can describe musical attributes in terms of parameters that specify *type* and *degree*, which we can use to augment the illuminative power of the aforementioned structural concepts. The type, degree and multiplicity ((dis)similarity and (dis)continuity can occur in relation to an individual parameter or multiple parameters simultaneously) of parametric change affects our overall sense of discontinuity. These parametric descriptions help us understand why some musical relations strike us as (dis)continuous and some do not.

*Degree* is of particular importance in such endeavors, for while structural descriptions typically require that we distinguish between *types* of content, they do not maintain the same sort of rigid standards for descriptions of *degree*. For example, consider the chorused piano / drums MOP highlighted in section 4.1.4 of Chapter IV.

We could use structure mapping to describe the relationship between the constituent entities so that our account would only consider the type of attributes involved: the drums in *entity A* are not chorused while the drums in *entity B* are. We could also describe this relation as a high-level transformation where *entity A* is transformed into *entity B* by introducing a chorus effect to the drums. Such a description would also only reflect the *type* of attributes involved in the relation. The problem is that neither of these descriptions gives us a sense of *degree*.

Indeed, *degree* significantly influences whether we perceive a relation as continuous or not. If a slight amount of chorus was added to *entity B*, the relation may not be particularly discontinuous. If a significant amount of chorus were added to *entity B*, we may perceive discontinuity. Not all types of specification are perceptually relevant. If one were to describe the transformation between *entity A* and *entity B* as in a 21% increase in the chorus mix, such a description may be technically accurate but isn't necessarily a reflection of our perceptual experience (maybe we don't notice a difference between 0 and 10%, maybe we are sensitive to changes in the 40-50% range, such as how we perceive frequency). The question then becomes how do we specify and represent *degree* in a perceptually-relevant way?

## 6.2. Dimensional representation

If we want to specify and represent *degree* in a perceptually relevant way, we need to understand how content (entities and attributes) corresponds to, and is defined by, higher-level reference *scales*. For example, in Planets' *O People* (Figure 6.1), we are struck by difference (and discontinuity in this case) when entity with the attribute *fast* is immediately followed by an entity with the attribute *slow* (an alignable difference).

[Click to listen](#)

**Figure 6.1. excerpt from Planets – *O People*.**

This differentiation requires that we have a sense of an abstract, higher-level reference to which the attributes of both entities correspond. This reference allows us to relate the individual attributes; thus, it is a *scale*. In the example, we have a sense of the abstract scale *tempo* to which both entities correspond via their attributes *fast* and *slow*. The scale, *tempo*, is common to, and thus an inherent part of, the attributes, *fast* and *slow* (attributes thus have a *type* and *degree* – tempo:fast and tempo:slow in this case). This commonality allows us to relate the attributes, and thus the entities, to each other. Thus, we are able to differentiate as we find commonality. It was stated before that our ability to perceive likeness requires our ability to perceive difference. We can now add that our ability to perceive difference requires our ability to perceive likeness (this is related to the earlier observation that people list *more* differences for similar pairs than for dissimilar pairs).


Once we are able to relate attributes to each other via a common reference, we can represent them on a continuum from which we can calculate a (dis)similarity quantity. The poles of this continuum are typically defined by a pair of opposites. We then orient attribute-percepts relative to those poles to arrive at point-locations that represent those attribute-percepts. The distance between the attribute-points will be a function of both individual *degrees* and the poles of the continuum. Both degree and the poles can be described in terms of generalizations such as *fast* and *slow*, or can be more precisely quantified, depending on the dimension. Once the poles and individual degrees have been specified, the distance between the attribute-points can be qualitatively or quantitatively measured. The resultant difference or (dis)similarity quantities can be used to augment our structural descriptions, which, in turn, can better illuminate our experience of musical (dis)similarity and (dis)continuity.<sup>44</sup>

To illustrate how this works, consider the previous example. We could define

---

<sup>44</sup> Importantly, this is not to say that all such points can be located in a (multi)dimensional space, that the distance between points is symmetric nor that dimensions are absolute and resistant to context or individual variation. Rather, such dimensions can be used in *specific contexts* in order to make the preceding general framework capable of more granular descriptions.


the poles as the slowest and fastest passages that a listener has experienced. The listener could then, in reference to this scale, locate each attribute-percept as a point on the continuum (see Figure 6.2).



**Figure 6.2. Tempo (dis)similarity continuum**

If we assign values throughout the continuum, we could measure the distance between points to derive a quantitative measure of similarity (the closer the points are to each other, the more similar are the attributes they represent). Given the scale set by the poles ( $a, b$ ), the quantitative value of (dis)similarity would be obtained by  $(d - c)$ .

We can then locate the (dis)similarity quantity on a difference (or likeness or similarity) continuum whose poles of *opposite* and *same* are connected by varying degrees of similarity (Figure 6.3). The quantity of (dis)similarity corresponds to qualitative judgments such as “same” (smallest quantity), “similar” (small quantity), “dissimilar” (large quantity) or “opposite” (largest quantity).



**Figure 6.3. Difference continuum.**

This (dis)similarity quantity can then be used to augment structural and transformational descriptions. In this example, the two entities have similar relational structures, and the attributes that contrast (*fast* and *slow*) play similar roles relative to those relational structures (both attributes articulate *tempo*). The attributes fast and slow

are therefore aligned, so we can call this relation an alignable difference. We can then further describe the relation as a tempo transformation (tempo describes parametric *type*); the *degree* of that transformation is specified by the (dis)similarity quantity measured. Putting all the pieces together, we can then specify our experience of discontinuity by articulating the structural characteristics, transformational type and the (dis)similarity quantity of this particular relation.

### 6.2.1. More complex dimensions


The above description of dimensional representation is somewhat simplified. That is, defining points on a dimension is not as simple as picking a pair of opposites, drawing a line between them and placing points on that line. As we will see shortly, dimensions become complex as we start to consider different types of poles and intermediaries. These more complex dimensional representations are useful to the extent that they more accurately reflect perceptual experience. Therefore, the following will explore some of the details of this complexity.

To address the complex possibilities of dimensional representation, Bianchi, Savardi and Kubovy (2011) first distinguish between types of opposites. They cite Ogden's differentiation between *series opposites*, the poles on a continuum, and *cut opposites*, where opposition begins immediately when a line is crossed (such as INSIDE-OUTSIDE). They liken this distinction to that between gradable adjectives and complementary opposites. Gradable adjectives (DIRTY-CLEAN) allow for degree distinctions between poles while complementary opposites are less flexible (MAN-WOMAN). Sometimes gradable opposites have a middle term (warm within COLD-HOT), sometimes they do not (DIRTY-CLEAN). Complementary opposites can also allow for some level of gradation (*further* INSIDE).

The authors also explore how a continuum (dimension) can be segmented, which is discussed within the category of "metric considerations." One possibility is that




continua are not symmetric. For example, in the case of SMALL-LARGE, it may be the case that we think “the scope of large things is greater than the scope of small things” (Bianchi et al., 2011, p.5). We can therefore divide a continuum (or dimension) proportionally according to that scope of each pole (top rectangle of each dimension in Figure 6.4).<sup>45</sup> We can further divide that continuum according to that which is *not* one pole or another (bottom rectangle of each dimension in Figure 6.4). Bianchi et al. implemented these ideas in an experiment, which yielded the following representations in Figure 6.4:



**Figure 6.4. Scope and division of poles. Bianchi et al., 2011, p.13.**

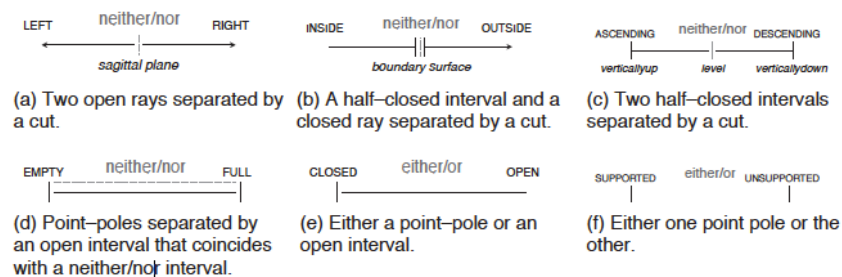
Bianchi et al. (2011) also identify “topological considerations”, which explore the possibilities of pole types and dimensional qualities. *Closed* poles refer to the minimum and maximum of a scale (EMPTY-FULL). *Open* poles are unbounded (SHORT-TALL). We can represent these ideas using two endpoints ( $a$ ,  $b$ ) joined by a line segment (the latter is referred to as an *interval*). “Each endpoint can be a pole or a cut” (3). Intervals can be closed ( $a$  and  $b$  are included in the interval), open ( $a$  and  $b$  are not included in the interval), half-closed (one endpoint is included and the other is not), a closed ray, an open ray (rays involve  $\infty$ ), or a point (Figure 6.5 depicts various types of intervals).

<sup>45</sup> The authors propose we derive these proportions by assigning phenomenal representations to “experience bins”, which we can then assign to one pole or another. The proportion of experience bins assigned to each pole determines where the continuum is segmented. See Bianchi et al. (2011) for a more in-depth discussion.



**Figure 6.5. Various intervals. Bianchi et. al, 2011, p.3.**

Between the poles can be an 1) intermediate state (neither / nor), which can be an 1a) isolated state or an 1b) open interval or 2) no intermediate state. Combining these elements allows for a variety of possible dimensions, some of which are illustrated in Figure 6.6.




**Figure 6.6. Various dimensions. Bianchi et al., 2011, p.4.**

To define a dimension then, Bianchi et. al ask 1) if a pole is a point or an interval, 2) if the pole is an interval, is it open or closed, 3) is there an intermediate and 4) if there is an intermediate, is it a point or an interval. Responses from these questions can be used to derive topological identity for any particular dimension. Topological identities can then be compared to determine the relative similarity of dimensions themselves. We can use topological qualities in isolation or in conjunction with metric qualities (as Bianchi et al. do) in order to characterize dimensions more comprehensively.

For the purposes of this project, the metric-topological character of dimensions is

important for two reasons. First, the way that we quantitatively or qualitatively characterize distance between points along any dimension, which we are interested in as earlier discussed, may be affected by the metric and/or topological identity of that dimension. Second, perceptual salience may vary depending on a dimension's metric and/or topological identity, that is, discontinuous relations that occur on certain kinds of dimensions may be more perceptually salient than others.

In regard to the first point, the (dis)similarity quantity is affected by how we define the poles of the reference continuum. Instead of the continuum described in Figure 6.2, imagine if the poles were defined as the slowest / fastest passages imaginable, rather than as experienced (Figure 6.7). Augmenting the range of possibilities extends the poles of the continuum and subsequently shrinks the distance between the two passages. As a result, the (dis)similarity quantity is smaller and the slow and fast passages would be judged as more similar (relative to Figure 6.2), despite the fact that the passages are phenomenally identical. This is something that we need to keep in mind when measuring distances.



**Figure 6.7. Tempo (dis)similarity continuum.**

The second point, that perceptual salience may vary depending on a dimension's metric and/or topological identity, is also important. Of course, the perceptual salience of a particular dimension will likely be the function of a number of elements including structural identity, context and individual variation, but it is nonetheless interesting to consider the possibility that some dimensions are inherently more or less perceptually salient. Such qualities would certainly be important in the analysis / prediction of

musical similarity percepts, so experiments that explore these possibilities could be interesting.

### 6.2.2. Identifying parametric dimensions

We must also consider how we can pinpoint the dimensions that are most relevant in musical (dis)continuity relations. We do not want to approach a (dis)continuity relation with a list of all the parametric dimensions along which its components can be perceived.<sup>46</sup> Instead, we want to focus on the most salient features, which we can then explore in greater depth using dimensional representation and comparison. As mentioned in chapter IV, asking participants (or ourselves as analysts and composers) about salient cues, musical boundaries or dimensions along which change occurs could be an effective way of choosing which parameters are relevant. When we have identified the most salient features, we must then take care to choose relevant dimensions and perceptually meaningful representational methodologies.

### 6.2.3. Parameters that resist dimensional representation

Some musical parameters, such as meter, rhythmic configuration and timbre, don't lend particularly well to the kind of dimensional description highlighted above. What sort of oppositional pair could be the basis for a timbre dimension? How could we describe the distance between the meters 3/4 and 7/16 in a musically relevant way? One approach to such problems involves deconstructing attributes into their constituents, which may be more amenable to dimensional representation. For example, we can describe timbre in terms of tonal / noise character, spectral envelope, amplitude envelope and frequency changes. We can then combine these lower-level attributes to form a single measure of the higher-level attribute, which can then be located on a

---

<sup>46</sup> Even if we were to have access to such a list, we would have to be careful to use it in a way that is cognizant of the problems with the featural approach discussed earlier: the most exhaustive of checklists will still be incomplete and perhaps even more importantly, such a method doesn't segment the features that are contextually relevant from those that are not.

(dis)similarity dimension.<sup>47</sup> Alternatively, we could depend more heavily on the transformational instructions necessary to convert one sound into another, that is, transformational distance would have a greater weight in some types of (dis)similarity descriptions. For example in the case of rhythmic configuration, transformational distance might be helpful in describing how close or far one rhythm is from another (and thus how (dis)similar they are). Alternatively, we can rely on measures that have been developed to represent these more convoluted parameters, such as Grey's multi-dimensional timbre space (1975). These sorts of possibilities reinforce the need to have a varied and flexible analytic approach. Any of our descriptive tools will be illuminative in some cases and less so in others. The ability to use varied tools will allow us to illuminate more of the richness of our (dis)continuous percepts.

To exemplify how the dimensional identification and representation process described above works in a musical context, I will consider an excerpt from Planets' *O People*. When we listen to this song, more than anything, we are struck by rhythmic change. Therefore, I will consider the most salient rhythmic parameters and associated dimensions that we can use to describe our experience of the music.<sup>48</sup>

### 6.3. Rhythmic parametric change in Planets' *O People*

Rhythmic change defines the discontinuity of Planets' *O People*:

---

<sup>47</sup> Of course, such a process would need to be sensitive to context and the levels of a hierarchy if it were to produce perceptually meaningful values.

<sup>48</sup> The following discussion will show how a larger musical category, such as musical temporal organization, comprises a variety of constituent dimensions that we can use to describe (dis)continuity relations. Since many musical entities contain pitch, timbral and durational attributes, we could similarly describe these primary categories in terms of their constituent dimensions. For example, in the realm of pitch, parameters could include pitch sets or registral distributions. Within the broader category of timbre, parameters could include orchestration, processing (filters, time-based effects), inharmonic transients at the beginning of tones, spectral flux (change of spectral envelope / harmonics over the course of a sound) and brightness (Grey, 1978; Bregman, 1990; McAdams, 1992). Dynamic changes can be measured via the perception of loudness and can apply to various parts of musical hierarchy, from individual events to the entire musical whole. In each of these fields, measurement models have been developed that can be applied to a (dis)continuity analysis.

[Click to listen](#)

**Figure 6.8. Rhythmic discontinuity in Planets' *O People*.**

The most salient characteristic of the first section (*section 1*) of this except is its energetic and flailing rhythmic identity. After a short transition, the subsequent section (*section 2*) adopts a strikingly different rhythmic character, inspiring a sense of discontinuity relative to the first section. We can describe the contrast that effects this discontinuity in terms of the following parameters: *number of elements*; *density*; *beats, meter and tempo*; and *synchronicity and simultaneity*. The format of this analysis will be to first describe the parameter and its dimension in an abstract way, which I will then apply to specific musical examples.<sup>49</sup> Because the categories *number of elements* and *density* are related, I will describe both of them first, after which I will discuss the other parameters.

### *Number of Elements*

If the duration of both entities in a relation is equivalent, contrast, and thus potentially discontinuity, can be created if the *number of elements* contained within each entity differ significantly (see Figure 6.9).



**Figure 6.9. Contrasting Number of Elements.**

Number of elements can be represented and compared on a dimension whose poles are defined by a point (none) and an open ray (many)<sup>50</sup>. The intermediary presents a somewhat complex problem. One could argue that the parts of this dimension closest to

<sup>49</sup> The dimensional examples that I will give are the product of my own judgments. Ideally, such questions could be examined in experimental settings to illuminate agreement and difference between the percepts of larger groups of people.

<sup>50</sup> The *many* pole could be closed if one were to define the dimension relative to the maximum number of elements.

the (none) pole should be represented with points, not a continuum, given they represent discrete values. Because we are interested in the perception of *number of events* and not just the physical facts, we can represent perceived number of elements as a continuum when we can no longer specify exact quantities. For example, if one entity contains 17 events and the next contains 28 events, we likely will not be able to describe exactly how many events each entity comprises, although we likely will be able to compare the number of events in a more general way. What this threshold is, exactly, depends on context and the characteristics of our perceptual / cognitive systems. In any case, *number of elements* is a hybrid dimension characterized by discrete points on one end and an open interval on the other. We could also compare the *perceived* number of elements to the *actual* number of elements to gain additional insight into how we interpret configurations in musical contexts.

### *Density*

Density is an important part of musical temporal perception and thus, contrast in density can create musical discontinuity. Often, we have little sense of what the actual durations are that separate notes yet we have a strong sense of how tightly they are bunched relative to the surrounding musical context. Density is a higher-level percept in that it is derived from number of notes per unit time, rhythmic configuration and rate. Density should not be confused with number of notes per unit time, as Figure 6.10 shows:



**Figure 6.10. Difference between number of notes per unit time and density.**

Both measures have two notes per unit period (measure) yet the density of notes differs between the two.

In terms of representation, as in the case of *number of elements*, we are primarily interested in describing perceived density rather than physical density. We can therefore ask listeners to locate perceived density associated with the entities that are adjacent to a discontinuity barrier on an appropriate dimension. This dimension would have open interval poles (sparse and dense are unbounded) and a neither / nor intermediary. We can then compare perceived density to actual density so that we can further understand how we perceptually interpret density in musical contexts.<sup>51</sup>

In *O People*, discontinuity is created through contrast along both of these parametric dimensions. The aforementioned energetic and flailing character of the first section is largely a function of the drum performance. This drum part strikes us as containing many elements and dense (even before hearing the subsequent sections). As section 2 settles into something more sane, which is characterized as containing fewer elements and less dense, our initial impression of the first section is confirmed. We thus could represent the discontinuity created along the dimensions of *number of elements* and *density* as in Figure 6.11:

---

<sup>51</sup> The question of representing physical density in a perceptually meaningful way is a complex topic. One of the main issues with describing physical density is choosing a perceptually-appropriate temporal period. One approach could be to center a temporal window at the point of contrast and extend the left window boundary until one event is reached and the right window boundary until two entities are reached. We can then represent the density on either side of the boundary as events / unit time. In Figure 6.10, the center would be the end of measure one / beginning of measure two, the left boundary would extend to the half note before the end of measure one and the right boundary to the second sixteenth note after the beginning of measure two. If the tempo of Figure 6.10 was quarter = 120, the density of *entity A* would be 1/1000 msec, while the density of *entity B* would be 1/125 msec. The advantage of such a description is that it is sensitive to the temporal regions closest to the point of discontinuity; those regions correspond to the most immediate and therefore salient experiences. The disadvantage of such a method is that it is sensitive neither to changing densities nor to the periods of time beyond the immediate area of the discontinuity boundary, which also likely influence our judgments. We thus may choose to extend the temporal window to include more events, but if we do this, we must consider how far we are going to extend the temporal window in each direction, and we must choose a method to describe densities that change over time. The methodologies that are best equipped to deal with such issues are the subject of future research.





**Figure 6.11. Number of elements (top) and density (bottom) contrast in *O People*. The blue circles represent section 1; the orange circles represent section 2.**

### *Beats, Meter and Tempo*

Beats, meter and tempo are closely related, and are among the most salient qualities that a musical passage can project; it follows that contrast within these parameters can create powerful discontinuities. The sense of an underlying pulse or beat is foundational to musical traditions across cultures. Our sense of beats is an emergent phenomenon that is produced from the relationships between individual sounds (it is emergent because such a percept does not inhere in any of the individual sounds). We typically organize beats into groupings (meter) that can occur at a certain rate (tempo).

To further understand how these musical qualities relate to each other and the musical surface, we can describe them relative to a hierarchy. In the case of beats, hierarchical level is both clear and opaque. It is clear that beats are higher-level than individual events (the former is derived from the latter). It is less clear whether or not beats are higher-level than rhythmic configuration. It seems most likely that our sense of beats, and thus a beat grid, is derived from rhythmic configuration (as well as from other musical characteristics such as accent and pitch patterns), and thus the former is higher-level than the latter. One could argue that beats are derived from the relations between individual sounds, as are rhythmic configurations, thus, rhythmic configuration and beat grids inhabit the same hierarchical level. The question is really: do we understand beats through configurations or configurations through beats? For

meter and tempo, we can say that both are higher-level than beats, given they are derived from the relationship between beats, although the relationship between meter, phenomenal sonic events and mental structures is complicated (see Palmer and Krumhansl, 1990 for a discussion of how we represent musical meter). Further, rhythms and meters can combine to form even higher-level organizations (see Cooper and Meyer, 1960).

Representing these qualities is straightforward in some instances and less so in others. In regard to the former, *tempo* can be located on a dimension whose poles are defined by two open rays (*slow* and *fast*), with a neither / nor intermediary. Alternatively, the poles could be closed if we defined them according to the extremes of the piece or larger set of pieces. Representing and comparing meters seems more difficult, for how similar are 4/4 and 5/16? One approach to represent meter is to separate the grouping component from the rate of beats component. Consider Figure 6.12:



**Figure 6.12. The same rhythm notated in two ways.**

These two rhythms are identical, even though one is noted with the time signature 5/4 and the other with 5/8. Thus, we are interested in the top number of a time signature, which indicates the grouping of beats, and the speed of those beats, which is a function of the bottom number and the indicated tempo. If we are comparing music that hasn't been notated in the Western convention, we need only concern ourselves with grouping size and rate of beats (we don't need to assign a beat to a somewhat arbitrary rhythmic value, such as a quarter note). We can represent grouping size with two open ray poles

(*small* and *large*) and a neither / nor intermediary.

In *O People*, contrast in meter and tempo contributes to the discontinuity featured in Figure 6.8. Metrically, the first section can be interpreted as functioning in 4/4 (it is indeed a loose 4/4, which is one of the qualities that gives the music its distinctive mercurial rhythmic character). More specifically, the music is organized into larger phrases that span two measures of 4/4 with the last two beats of the second measure consisting of a quarter note triplet (the triplet is emphasized to such a degree that one could argue that it is a significant, but brief, rate change in itself, rather than as a part of the 4/4 measure). The tempo of this section is roughly  $q = 158$ . Conversely, the subsequent section alternates between 7/8 and 5/8 at a slower  $q = 141$ . In this case, 7/8 and 5/8 are notational distinctions, rather than reflections of changes in the underlying pulse relative to the 4/4 section, thus, 3.5/4 and 2.5/4 are more appropriate descriptions that better align with our perceptual experience of the music.

We can compare these contrasts in a number of ways. First, we could describe the meter of the first section as unchanging, which contrasts with the meter(s) of section two, which changes every measure. We would describe this attribute as a feature, and not a parametric dimension, according to the definitions in chapter I. Next, we can represent the discontinuity's tempo and meter contrasts on the following parametric dimensions (Figure 6.13):



**Figure 6.13. Tempo (top) and metric grouping size (bottom) contrast in *O People*. The blue circles represent section 1. The orange circle in the tempo dimension represents section 2. The orange circle in the metric dimension represents 7/8, the yellow circle represents 5/8.**

The level of description in which we are interested will dictate which metric representation we use. If we are interested in the highest-level description that spans the greatest temporal period, we can characterize the discontinuity as a juxtaposition of metric stasis and metric change (the featural description mentioned above). If we want to take a lower-level perspective and consider the character of the individual meters, we may utilize a representation such as the bottom dimension in figure above, which gives us some sense of the different kinds of grouping sizes involved in the larger relation. If we zoomed even more closely to the point of discontinuity, we may only want to look at the groupings on either side of the boundary, which would include  $3.5/4$  ( $7/8$ ) and  $4/4$  (the orange and blue circles in the figure above). All of these levels are in play in a listening experience, and all contribute to our sense of discontinuity.


### *Simultaneity and Synchronization*

When a musical texture contains multiple voices, we can describe the rhythmic relationship between those voices in terms of simultaneity and synchronization. In a general sense, simultaneity and synchronization refer to how different voices interact with each other relative to a temporal reference. If *voice A* and *voice B* both occur at *time 1*, we can describe them as simultaneous. If *voice A* and *voice B* are related by a particular temporal function or scale, then we can describe them as synchronized. When a simultaneous or synchronized section is contrasted with non-simultaneous or non-synchronized section, discontinuity can be created.

When listening to a particular musical selection, we are able to perceive and subsequently represent a level of simultaneity that characterizes the section as a whole. Our sense of simultaneity in this context is a function of the number and salience of alignments that occur between voices during that selection. We can ask listeners to represent this overall sense of simultaneity with a value between 0% and 100%. This

value can then be located on a dimension with closed point poles (0% and 100%) and an open interval neither / nor intermediary. While we usually articulate such perceptions with ease, the phenomenal context that produces such perceptions is not always as clear. That is, we often aren't sure of exactly how many musical events aligned with each other, or which alignments were most salient, even though we are able to give a general sense of the result of such alignments. This ambiguity inspires us to explore the phenomenal context more comprehensively so that we can learn more about our perceptive tendencies.

To better understand the phenomenal context that our overall simultaneity judgments reflect, we can incorporate physical descriptions of aggregate simultaneity. One way to do this is to describe the overall simultaneity of a musical selection as a function of the number of coincident note events relative to the total events for each voice. In Figure 6.14, we could say that 20% of the top voice is simultaneous with the bottom while 50% of the bottom voice is simultaneous with the top voice. We could average these values to derive an overall simultaneity value of 35%. We could also describe simultaneity relative to the composite rhythm. There are two moments of simultaneity out of seven rhythmic iterations, which equates to a simultaneity value of ~30%.



**Figure 6.14. Simultaneity between two voices.**

We could compare these physical descriptions to perceptual ratings to infer what

phenomenal characteristics listeners perceive and respond to when judging the overall simultaneity of a musical passage.<sup>52</sup>


The idea of synchronization is related to simultaneity, but differs in several important ways. First, two different voices could be perceived as synchronized if they both adhere to a higher-level structure such as a beat grid, even if the phenomenal constituents of those voices don't occur simultaneously. We also must consider the possibility of different kinds of synchronization (which are perhaps better described as types of coordination), such as phase offsets. If two voices are characterized by identical rhythms, but are offset by some value so that the voices don't share a beat grid, we still have a strong sense that the temporal locations of each event in one voice are precisely related to the temporal locations of each event in the other. In this way, we can understand the voices as synchronized as the latter essentially involves alignment: here the voices are aligned, one is just offset temporally from the other. In terms of representation, if contrast comes from one set of entities being synchronized to a beat grid and another not, we can describe such contrast as a feature. If one set of entities are synchronized / phase-offset and another are offset by a different duration, we can represent the degree of offset on a dimension. If one set of entities are synchronized / phase-offset and another are not synchronized at all, we can also use a two-value representation that reflects those properties.

In *O People*, section one and two differ significantly in terms of the simultaneity of their instrumental components, which contributes to the discontinuity between the sections. In section 1, the temporal articulations of the drums are largely independent from that of the guitar; the two only occasionally converge to accent shared rhythmic structure. One may argue that because of the density of the drum part, there is, in a

---

<sup>52</sup> There are several other considerations that we must address in this context. One is that it is likely the case that we perceptually weight values depending on metric position. That is, simultaneity on beat one of a measure would be felt as stronger than simultaneity on beat four. We also must entertain the possibility that one voice may be more perceptually prominent than the other, so that A's simultaneity with B is more salient than B's simultaneity with A. These factors may affect our perception of the overall simultaneity of a musical selection.

physical sense, a significant amount of simultaneity between the drums and guitar. This kind of simultaneity is of less interest than the kind that reinforces or amplifies the connection between the instruments or some structure to which both instruments relate, such as a beat grid. We can refer to the former as *overlap* and the latter as *simultaneity*. Therefore, in section 1, the drums primarily overlap with the guitar part. In section two, the drums articulate the metric structure of the guitar part to a much greater degree, projecting a number of structurally-meaningful simultaneities. We can represent the discontinuity's simultaneity contrast on a dimension as follows (Figure 6.15):



**Figure 6.15. Simultaneity between the guitar and drums in *O People*. The blue circle represents section 1, the orange circle represents section 2.**

### *Combining parameters*

Combining the above rhythmic dimensional representations (number of events / density; beats meter and tempo; and simultaneity and synchronization) allow us to paint a more detailed picture of the *degree* of rhythmic contrast that we perceive in Planets' *O People*.<sup>53</sup> To bring this picture to fruition, we first can measure the distance between points in each of the dimensions above to derive values that represent (dis)similarity. Next, we can consider the relative perceptual salience of each dimension in the musical context being considered, which can result in a weighting of each value. We can then synthesize these values to derive a measure of (dis)continuity for a given musical moment. We can repeat this process for other discontinuous moments in a work to illustrate how discontinuity functions in larger forms. Such an approach enhances our

<sup>53</sup> Note, I chose these particular dimensions as they reflect the most salient types of change in this particular excerpt. We could enhance our description of this discontinuity with more dimensional representations, such as the kinds detailed in Appendix A. Further, as previously noted, different musical examples may project different musical parameters.

descriptions of relations, whether structurally- or transformationally-based, in a contextually-sensitive way, which will allow us to better understand why a musical juxtaposition strikes us as discontinuous. As appropriate dimensional representation is key to this effort, it follows that further experimentation is necessary to define perceptually-relevant musical dimensions. Such research will enable us to conduct more comprehensive analyses of musical (dis)continuity.

## **6.4. Interaction of Parameters**

The above discusses musical parameters in isolation but in most music, changes can and do occur along multiple parameters simultaneously. As one would expect, the topic becomes exceedingly more complicated when multiple parameters are simultaneously involved. Multiple parameters can interact, via the entities that project those parametric qualities, in a variety of ways depending on parametric type and degree, temporal relationships and context. In regard to how types of interactions, parameters can compete for perceptual saliency, combine to form higher-level entities, combine to form higher-level parameters and create perceptual segmentation boundaries.

### **6.4.1. Parameters competing for perceptual saliency**

When multiple parametric attributes of entities are present, they compete for perceptual saliency. For example, we may notice that flute melody changes, but not that the timbre of the flute also changes, even though both may occur simultaneously. One way that we can examine the relative saliency of different parameters is to consider the extent to which they promote perceptual groups. If rhythmic configuration is more group-promoting than registral location, then music that maintains rhythmic configuration while changing registral location may be perceived as more continuous than that found in the opposite scenario. We can also use the idea of the systematicity



principle, as described in chapter IV, to describe which parameters gain perceptual salience. Such interactions will depend on the type of parameters, the degree of change within those parameters and the context within which they occur. Having an idea of how parameters compete for perceptual saliency can help us weight parametric representations appropriately when multiple parameters are simultaneously present.

#### **6.4.2. Parameters combining to form higher-level entities**

Multiple parameters can also combine to form higher-level entities. An example is a theme, which usually contains the parameters rhythm, pitch, timbre and dynamics. When we try to describe thematic discontinuity then, we can do so in the context of any of these parameters and the dimensions that represent them. Because several parameters can characterize a theme, we must identify which are more salient and which are less so, given the context. It may be the case that a theme is defined primarily by its rhythmic configuration or its pitch sequence. After we identify which components of a theme are perceptually most important, we can appropriately weight their corresponding dimensional representations.

#### **6.4.3. Multi-level Parameters**

Individual parameters can also combine to produce more comprehensive parameters that can describe entities on multiple hierarchical levels. These parameters, which can be mapped onto dimensions (poles indicated parenthetically), include *energy*, a function of amplitude, density and durations (low – high); *topography*, understood as a degree of “standing out”, which can be applied to register, amplitude, timbre, rhythmic configuration or the musical whole (flat – uneven); *repetition* (none – identity); implied *motion*, which involves expectation and directionality (stasis – perpetual) and *contour*, which can apply to specific parameters like pitch or dynamic sequence or the musical whole (different measures of contour have been developed according to different

parameters, such as melodic contour). Each of these comprehensive parameters inhabits an elevated part of a hierarchy given each results from the combination of numerous low-level parameters. Given our tendency to favor higher hierarchical levels, it is not surprising that we often reference such categories when articulating our experience of musical (dis)continuity.

#### **6.4.4. Parameters creating perceptual structural segmentations**

When change-inspired boundaries occur across multiple parameters, the perception of higher-level structural segmentations can be produced. For example, if the boundary of a melody and a rhythmic configuration occur simultaneously, we may perceive the start of a new section, “section” being a higher-level organization than both melody and rhythm. The intensity of a combined segmentation boundary thus is a function of the number and strength of constituent parametric grouping boundaries. In *Music and Memory*, Bob Snyder refers to this phenomenon as *intensification*. He states that as long as change along each parameter is of a sufficient quantity, such change can create boundaries that are higher-level than those of the individual constituents [204]. This assertion is supported by the work of Deliege and her colleagues, as mentioned in chapter III, which states that the perceptual salience of segmentation boundaries is a function of the degree of change that occurs between the entities adjacent to such boundaries.

While we can agree that change corresponds to perceptual segmentation boundaries, it is not necessarily clear what kind of change or how much change is necessary to create such percepts. Snyder doubts that change in one parameter would be enough to demarcate a section. Changes in two parameters may be possible, although it would depend on what parameters change and by how much. Snyder claims that change along three or more parameters is likely strong enough to indicate sectional change, and that there is some research supporting this statement (he cites

Miller and Johnson-Laird, 1976) (204). While these seem like intuitively reasonable guidelines, we should take care not to accept them too rigidly, for there are a number of factors beyond just type of parametric change that affect the perception of segmentation boundaries.

Indeed, the strength and hierarchical level of a perceptual boundary is influenced by a number of factors; of these, context is one of the most important (Snyder acknowledges this point). For example, in the first movement of Carter's *Sonata for Flute, Oboe, Cello and Harpsichord*, we perceive change as voiced through the harpsichord part in ma. 23 (see Figure 6.16). Analyzing this moment in detail, we can describe this change along several parameters. First, in terms of rhythmic configuration, mm. 21-22 (and the very beginning of the piece) is characterized by inter and intra-voice heterogeneity. Within a voice, this means quick rate changes, such as the flute in ma. 22, which plays triplet eighth notes for two beats and then shifts to sixteenth notes for the next two beats (the oboe does the reverse in ma. 22, playing sixteenth notes and then triplet eighth notes). Between voices, different superimposition rates often sound simultaneously (usually three against four, as in ma. 22). In contrast, the harpsichord in mm. 23-34 plays a string of isochronous sixteenth notes that lasts almost two measures. There is also dynamic change: mm. 21-22 (as well as the preceding measures) are characterized by rapid dynamic oscillation, a function of short crescendi and decrescendi in the winds and cello, while, conversely, the harpsichord in mm. 23-24 is unwaveringly forceful. The temporal density of the harpsichord part also changes at mm. 23: five rhythmic values fill the ma. 22 versus fourteen in ma. 23. Finally, the harmonic density of the harpsichord changes at ma. 23. Mm. 21-22 are generally characterized by sparseness, a function of individual notes and chords that are registrally dispersed. In ma. 23, the harpsichord alternates between more tightly packed chords that typically don't span more than a fifth (with a few exceptions in the left hand).

This example illustrates how context, and not just absolute change, influences

discontinuity judgments. Even though I described changes along four different parameters (and it is possible to describe more changes), we don't necessarily experience the point of change as a strong sectional boundary or a moment of discontinuity. Part of the reason for this is because the changes are voiced for two measures; had they continued throughout a more extensive part of the piece, a segmentation boundary might have been retroactively inspired. Additionally, this percept is influenced by the surrounding musical context, which up to this point in the piece can be described as a collection of diverse interjections that convey a range of tempi, dynamics, intensities and contours. Such a heterogeneous past expands the possibilities of future directions. Thus, we find the harpsichord's changes at ma. 23 a point of contrast, but we do not find it so significant as to establish a structural segmentation; such change seems within the realm of possibilities as defined by the previous music. Had the previous music been less diverse, then these changes might have been enough to inspire a perceptual boundary (Snyder also mentions this possibility [204]). In any case, this example shows how change in at least four parameters is not enough to effect a segmentation boundary, in contrast to the general guidelines articulated by Snyder.

The musical score consists of four staves. The top staff (Flute) begins with a triplet of eighth notes marked *mf*, followed by a phrase marked *f*. The second staff (Oboe) starts with a triplet marked *ff*, then a phrase marked *mf*, followed by a phrase marked *p sub.* and another marked *mf*, ending with a triplet marked *ff*. The third staff (Cello) begins with a triplet marked *ff*, then a phrase marked *f marc.*, followed by a phrase marked *mf* and ending with a phrase marked *f*. The bottom staff (Harpsichord) features a *Tutti: (ff)* marking and a *marc.* marking. The score includes various musical notations such as triplets, slurs, and dynamic markings.

Figure 6.16. Mm. 21-24 of the Carter's *Sonata for Flute, Oboe, Cello and Harpsichord*, mvmt I.

[Click to listen](#)

**Figure 6.17. Carter's *Sonata for Flute, Oboe, Cello and Harpsichord*, mvmt I mm. 1-24. Perf. Contemporary Chamber Ensemble.**

John Zorn's *Road Runner* also illustrates how multi-parametric change and contextual factors influence perceptual segmentation boundaries. In this piece, frequent high intensity changes demarcate short sections. From one perspective, the piece verifies Snyder's claims as multi-parametric change effect perceptual segmentation boundaries. From another perspective, as in the Carter example, the piece shows us how context can affect how many parameters must change to demarcate a new section. In *Road Runner*, we learn that a structural boundary requires simultaneous change along many parameters. Given these requirements are so extreme, it may be that change involving two or even three parameters may not be enough to establish a new boundary (such could be the subject of a future experiment). Both the Carter and Zorn examples show that context has significant influence on the way that we perceive segmentations as a function of parametric change.

In addition to context, there are other important factors involved in whether parametric change will be perceived as a segmentation boundary, which include the degree of change and the kind of transition between the entities. In regard to the former, a tempo change from quarter = 100 to quarter = 105 might not be enough to effect a segmentation but a change from 70 to 150 might be. In regard to transitions, a short silence separating parametric change is more likely to be considered a segmentation boundary than a transition where one parameter gradually morphs into another over the course of a minute.

It follows that we should avoid rigid guidelines regarding how parametric change affects perceptual segmentation boundaries and instead consider the number, type and degree of parametric changes, the transitions that separate those changes, and the musical context within which these various elements interact. Because structural

segmentations significantly influence our perception of discontinuity, exploring how parameters combine to form grouping boundaries is a potentially fruitful methodological approach for analysis or composition.

#### **6.4.5. Research on inter-parametric relationships**

While inter-parametric psychological studies are comparatively limited versus their intra-parameter counterparts, there are some examples in the literature. Firmino and Bueno examined the relationship between pitch space distances (via chord progression modulations) and temporal perception (2009). Moelants and Van Noorden investigated how pitch interval separation affects polyrhythm perception (2005). Krumhansl considered the relationship between pitch and rhythm in the context of melodic organization and phrase structure (2000). Multi-parametric interaction is an area that requires further experimental attention given its presence is ubiquitous in real music.

### **6.5. Associations**

Wholistic and component properties, in isolation or in combination, can evoke mental associations. We can relate, and thus create (dis)continuity between, entities based on the mental associations that they inspire. Mental associations can connect an entity to other musical elements such as works, physical instruments, particular performers, performance style, or to non-musical elements such as clothing conventions, specific experiences, physical locations or historical time periods. While entities on any level of a grouping hierarchy can inspire associations (both *wholistic properties*, such as those connected with a song, and *component properties*, such as those connected with a flute timbre, can inspire associations), there seems to be a correspondence between the hierarchical level of a property and the hierarchical level of the associations it inspires. For example, an entire song is more likely to trigger broad cultural associations than a

single drum sound (both ‘song’ and ‘broad cultural qualities’ are higher-level entities), but this is not always the case: a snare drum could trigger a rich tapestry of abstract cultural associations, depending on the snare drum’s distinctness, conventional use and the listener’s experience. The kind of association that is inspired depends on context.

Of the myriad of associational dimensions along which one can create contrast, style (particularly as style represents genre) is one that is frequently explored in discontinuous compositions. Stylistic juxtapositions are capable of creating dramatic contrasts given the diversity of musical styles in the world. John Zorn and his group Naked City’s *Snagglepuss* is a prime example of polystylistic juxtaposition: free jazz, straight jazz, funk, rock and noise genres (among others) are feverishly traversed in two minutes and twenty seconds. Alfred Schnittke’s *Concerto Grosso No. 1* also exudes stylistic discontinuity, described by the composer as a "a play of three spheres, the Baroque, the Modern and the banal" (Kim). Girl Talk’s *Pump It Up* comprises a diverse set of sounds that are associated with a wide range of musical styles including

How do we describe such associational juxtapositions? One strategy could consider change along the musical parameters that inspire such associations. A move from rock to jazz can be described in terms of instrumentation, chord voicing and rhythmic patterns, where contrasts are measured on the appropriate dimensions. Alternatively, we can discuss stylistic affect holistically by invoking metaphors like “smooth” or “heady” and / or by identifying non-musical associations such as performance venues or clothing conventions.

## 6.6. Approaching Parametrically-Rich Music

Because music is parametrically rich, we must form a strategy that will guide us towards the parameters that are most perceptually salient. As previously mentioned, revealing every possible dimension along which music can be described is not an efficient approach, given such a feature list would be prohibitively long, if not infinite.



Instead, we must identify parameters and dimensions in a contextually relevant way, which we can do by asking participants about salient segmentation boundaries (Deliege) or by asking them to identify relevant parameters directly. To illustrate the latter, Bianchi, Savardi & Cattazzo (2002) presented participants (amateur jazz musicians) with short jazz excerpts and asked them to focus on any kind of change that they heard. Participants then had to identify the dimensions, defined by two contrary properties, within which those changes developed (Savardi, 2008) (Figure 6.18). This is a promising model as it ensures that the parameters / dimensions are perceptually relevant and not arbitrarily established. Such a method produces a starting point and not a destination, as the list in Figure 6.18 is still likely too long and needs to be filtered according to those parameters that are most perceptually salient in moments of (dis)continuity.<sup>54</sup> Such filtering could occur through further experiments or musical intuition. Once we have identified the most relevant parameters, we can then represent them on appropriate dimensions, as previously described. The resultant description will be perceptually-informed, musically meaningful, thorough, and will have been derived from tractable methods.

---

<sup>54</sup> The table gives us some sense of how many dimensions are perceptually in play during a listening experience (note how many of these dimensions haven't been mentioned to this point). The number of dimensions here supports the earlier position that a comprehensive feature list for any musical example is prohibitively long.

Excerpt	Number of dimensions	Dimensions and frequency of recognition
1	11	open–close (7); repetition–crescendo (7); together–solo (6); change–repetitiveness (6); downbeat–upbeat(6); foreground–middle ground (6); medium low–medium high (5); movement–stillness (5); in synchrony–out of synchrony (5); same volume–different volume (5); movement within movement –movement within stillness (5).
2	13	solo–accompanied (6); unison–quarta (6); single note–appoggiatura (6); melodious–harmonious (6); foreground–background (5); repetition–crescendo (5); wide intervals–close intervals (5); self-sufficient–dependent (5); legato–fragmented (5); high–low (5); multi-themed–mono-themed (5); singular–plural (5); downbeat–upbeat(5); accidental–jumps (5).
3	5	fusion–segregation (7); solo–accompanied (6); parallelism–divergence (6); exchange–maintain (5); downbeat–upbeat(5).
4	9	same volume–different volume (7); multi-themed–mono-themed (6) open–close (6); slow–fast (6); gradual start–abrupt start (6); change–repetitiveness (5); complementing–superimposing (5); repetition–crescendo (5); rising–falling (5).
5	12	solo–accompanied (8); change–repetitiveness (7); unexpected–foreseen (7); slow–fast (6); accidental–jumps (5); melodious–harmonious (5); rich phrasing–poor phrasing(5); rising–falling (5); repetition–crescendo (5); con brio–pedantic (5); high–low (5); same volume–different volume (5).
6	4	solo–accompanied (9); foreground–middle ground (7); same volume–different volume (6); different instruments–same instruments (5).
7	10	solo–accompanied (7); rich phrasing–poor phrasing (7); same volume–different volume (7); high–low (6), movement–stillness (6); tension–relaxation (6); foreground–middle ground (5); movement within movement –movement within stillness (5); complementing–superimposing (5); melodious–harmonious (5).
8	14	same volume–different volume (7); tension–peace (7); repetition–crescendo (7); tension–relaxation (6); downbeat–upbeat(6); fusion–segregation (6); multi-themed–mono-themed (6); change–repetitiveness (5); together–solo (5); in synchrony–out of synchrony (5); open–close (5); rising–falling (5); melodious–harmonious (5); different instruments–same instruments (5).

Figure 6.18. Identified Dimensions. The number of participants recognizing the dimension is indicated in parentheses. Savardi (2008), p. 165.

## 6.7. Summary

The previous chapters discussed structural characteristics of (dis)continuity and (dis)similarity perception. In order to illustrate these structural characteristics more comprehensively, we can describe musical entities, attributes and relations in terms of

parameters, which we can represent on appropriate dimensions. We can define dimensions according to a pair of contraries using a variety of poles and intermediaries. We can then represent musical attributes on these dimensions, which will allow us to measure relative (dis)similarity and (dis)continuity. We can then weight these representations according to their contextual salience, after which we can combine them to derive an overall measure of discontinuity. Applied to music, we want to only consider those parameters that are most perceptually salient, which we can identify by asking listeners to note segmentation boundaries or discontinuous moments. A discontinuity relation from Planets' *O People* is used to illustrate how this process works.

Parametric interaction influences how we perceive (dis)continuity and structure in music. In regard to interactions, parameters can compete for perceptual saliency, combine to form higher-level entities, combine to form higher-level parameters and create perceptual segmentation boundaries. The latter is particularly important in the case of musical discontinuity as notions of discontinuity are often produced by segmentation boundaries. Generally, the more parameters change, the more intense the segmentation boundary will be, but other factors such as context, degree of change and transitions also play important roles in the perception of a segmentation boundary.

Parameters, via the entities that correspond to them, can also evoke associations with works, performers, styles, physical locations, etc. As entities relate to each other via associations, association becomes a vector along which we can create (dis)continuity. One kind of associational contrast, stylistic contrast, is frequently explored in discontinuous compositions. We can describe these associational contrasts in terms of constituent lower-level parameters, or we can use concepts that reflect qualities of the wholes of the associations.

We must also address how we choose the most relevant parameters when interpreting or creating discontinuity. I propose that salient parameters be specified through experiments that ask participants to either identify segmentation boundaries or

identify the parameters along which the music changes. From there, we can choose the representational and comparative methods that most effectively reflect our experience.

## 7. Chapter VII: Musical Works

I will synthesize the ideas of the preceding chapters in the context of four of my compositions: *for steps that grow when climbed*, *Breeding in Pieces*, *Push for Position* and *From Here to There*. As discussed previously, I will not use a uniform, inflexible method to approach these pieces. Instead, I will use the explanatory methods that best suit the particular musical situation. In some cases, structure mapping illuminates the nature of the (dis)continuities, in others, transformational distance is most useful. In some, I will consider how a listener perceives discontinuities (or structural segmentation boundaries), in others, I will consider how (dis)continuity was constructed from a compositional perspective. This approach will allow me to discuss various kinds of (dis)continuities and how those (dis)continuities function in diverse musical contexts.

### **for steps that grow when climbed**

2008-2009, revised 2011

for flute, oboe, clarinet, bassoon, tom-toms, marimba, piano, viola, double bass and electronics.


### *Perceptual Discontinuity and Segmentation*

To highlight the work's most salient discontinuities, I employed a methodology inspired by Deliege that instructs participants (here, the author)<sup>55</sup> to identify and quantify discontinuous moments during the course of playback. To enable this, I wrote a software program that prompts participants to press the enter key on the keypad when

---

<sup>55</sup> The depictions and descriptions in this chapter reflect the author's percepts. As the methodology becomes more developed, the goal is to involve a greater number of participants in formal experiments to see how discontinuity is perceived across groups of different people.

they experience a musical discontinuity, after which a dialog box appears requesting a value (0-9) that represents the strength of the discontinuity. Each response is timestamped relative to the audio playback when the enter key is pressed, not when the discontinuity value is entered, in order to accurately quantify reaction time. These results are shown in Figure 7.1 and Figure 7.2:



**Figure 7.1. Perceived Discontinuity in *for steps that grow when climbed*. The x-axis represents time within the piece (seconds) and the y-axis represents the strength of perceived discontinuity.**

00:17	4
00:28	4
00:43	3
01:39	7
01:50	5
02:07	5
02:12	3
02:22	2
02:35	4
03:33	9
03:53	7
04:07	4
04:46	5
06:10	5
06:15	4
06:25	7

**Figure 7.2. Table of discontinuity judgments in *for steps that grow when climbed*. The first column represents time within the piece and the second column represents strength of the perceived discontinuity (scale 0-9).**

We can analyze these moments more closely to identify the parameters along which change occurs. We then can represent these changes on appropriate dimensions to describe why certain moments appear as more discontinuous than others.

To begin, let us consider the most intense discontinuity perceived, which occurs at time 3:33 and was rated “9”. The strength of the discontinuity at 3:33 is a function of transitional type and change along a number of parameters including dynamics, acoustic / electronic, structure and beat-based / non-beat-based. Indeed, the work’s defining characteristics and most salient discontinuities typically involve these parameters, so looking at this moment in detail will help us understand the piece more generally. As mentioned in the discussion of parameters, we can describe a discontinuity as a segmentation that creates two *entities*: *A* and *B*. In the case of the aforementioned discontinuity, *A* ends at 3:32.27 and *B* begins at 3:32.75, with the two sections separated by a 478 msec silence.

### *Transition*

The transition in this relation is characterized by silence. The length of the silence in this case, almost half a second, contributes to the perceptual strength of the segmentation. Two other significant silences (defined as -50 dB and lasting at least 200 msec) occur in the piece: at 1:38.4, duration 670 msec and at 1:48.4, duration 210 msec. All of these silences are associated with perceived discontinuity. Referring to Figure 7.2, a discontinuity of strength 7 was identified at 1:39 and a discontinuity of strength 5 was identified at 1:50. As previously noted in chapter III, silence is one of the simplest and most powerful ways to create discontinuity.

### *Dynamics*

Dynamic contrast is a fairly regular characteristic of the piece, and the most extreme changes have important influence on discontinuity percepts. In the case of the contrast at 3:32.75, the RMS volume of the five seconds preceding the end of A is -17.4 dB, while the five seconds following the beginning of B is -31.6 dB; a difference of 14.2 dB. We can compare this to the other most significant dynamic contrasts in the piece to gain some sense of the influence of dynamic change on discontinuity perception (Figure 7.3).

3:32.75	14.2 dB	3:33	9
3:52.890	13.4 dB	3:53	7
4:44.659	18.1 dB	4:46	5
6:09.086	17.9 dB	6:10	5
6:24.290	13.2 dB	6:25	7

**Figure 7.3. Dynamic contrast in *for steps that grow when climbed*. Columns from left to right: actual time of contrast, volume difference (RMS), time of perceived contrast, discontinuity rating.**

Figure 7.3 shows that every one of the most significant dynamic contrasts in the piece was accompanied by a percept of discontinuity. At the same time, there is a wide range of discontinuity ratings associated with these moments of dynamic contrast, thus, dynamic contrast is not the only discontinuity-promoting factor that influences our percepts. In fact, the most significant moments of dynamic contrast, heard at 4:44 and 6:09, had some of the relatively lowest discontinuity ratings. It thus seems that dynamic contrast is only one factor that acts in combination with other factors to create a sense of discontinuity in this context.

### *Acoustic – Electronic dimension*



The transition between acoustic and electronic textures is one of the defining features of the piece; when such transitions meet the right conditions, discontinuity is created. In Western culture, we have become sensitive to the acoustic / electronic dimension given its prevalence in contemporary music, and these qualities now exert significant influence on how we identify and define musical voices. Novello et al. (2011) identify the acoustic-synthetic dimension as one of the primary factors in similarity judgments (along with slow-fast and vocal-non-vocal). In *for steps that grow when climbed*, the acoustic entities are voiced through the recording of the instrumental ensemble including flute, oboe, clarinet, bassoon, tom-toms, marimba, piano, viola and double bass. The electronic entities are voiced through diverse means, ranging from processing of acoustic sounds to textures created via software synthesizers. The acoustic and electronic trajectories are juxtaposed, synthesized and combined in various ways. The juxtaposition of these trajectories creates salient discontinuities.


The discontinuity at 3:32.75 contains one of the most dramatic contrasts within the acoustic-electronic dimension, which we can represent using the methods described in Chapter IV. First, we must determine the nature of the poles and intervals. In this case, we can consider ACOUSTIC and ELECTRONIC as closed point poles separated by an open neither / nor interval. We can locate B at the acoustic pole given there are no electronic elements during that passage.<sup>56</sup> The dimensional location of A is a bit more complicated. Within the context of the piece, A is electronic extreme but it is still clearly derived from acoustic sounds, so one could argue that we locate A near the electronic pole, as in Figure 7.4.



**Figure 7.4. Representation of 3:32.75 on an Acoustic-Electronic dimension defined by the context of the piece. The blue circle represents A and the orange circle represents B.**


<sup>56</sup> Other than the fact that B is obviously a recording that very much includes the influence of electronics.

Alternatively, if we define the electronic pole in the context of all electronic sounds, the location of B would be closer to the acoustic pole, given we can imagine a range of sounds that are further distanced from acoustic associations than B (see Figure 7.5).



**Figure 7.5.** Representation of 3:32.75 on an Acoustic-Electronic dimension defined by the context of all electronic sounds. The blue circle represents A and the orange circle represents B.

Depending on which contextual perspective the listener adopts, different levels of discontinuity will be perceived. In this case, it intuitively seems that Figure 7.4 is a more accurate representation of our experience, for the perceptual discontinuity is extreme both in terms of the acoustic-electronic dimension and the piece overall (supported by the 9 ranking). Accepting this representation, we can depict the changes in acoustic-electronic dimension over time, such as in Figure 7.6.



**Figure 7.6.** From top to bottom: timeline, perceived discontinuity and acoustic-electronic dimension in *for steps that grow when climbed*.

The representation of the acoustic-electronic dimension (Figure 7.6) tells us something about how the degree and duration of change affects our sense of

discontinuity. Most shifts along the acoustic-electronic dimension correspond to discontinuity judgments (at times 0:17, 1:39, 1:50, etc.) as long as they occur within relatively brief amounts of time. When shifts are separated by longer transitions, such as between 0:43 and 1:00, discontinuity is *not* always perceived, even though the entities on either side of that transition could be quite different in character. For example, there is a dramatic shift from 0:43 to 1:00 along the acoustic-electric dimension that is not accompanied by a discontinuity percept. The lack of perceived discontinuity is not because of a lack of contrast between the entities on either side of the transition, for if we juxtaposed the material surrounding times 0:43 and 1:00 without any intermediary transition, the result would sound like Figure 7.7, which would likely promote a sense of discontinuity:

[Click to listen](#)

**Figure 7.7. Juxtaposition without transition in for steps that grow when climbed.**

In the context of the work, the discontinuity-promoting power of this contrast difference is weakened by the content and temporal qualities of the transition. This example reinforces the discussion in chapter III that describes how (dis)continuity is affected by entity contrast *and* transitional attributes.

#### *Structural distinctions*

In *for steps that grow when climbed*, entities whose structure is clear are contrasted with entities whose structure is not, creating powerful discontinuities. As described in chapter IV, structural representation involves identifying the features and function of components relative to each other and the contextual environment as a whole. The roles that instruments or voices play is clear when they articulate form, define rhythmic structure, establish harmonic territory, etc. There are also musical textures that resist

these kind of structural distinctions in which individual relationships and roles are not easy to distinguish. For example, in Xenakis' *Concret PH*, we have some sense of individual components, but they come in vast numbers, are ephemeral and are highly similar. Such a texture resists the kind of structural representation described in the case of *My Reverie* or *Ordinary World*.

Because we are interested in describing the qualities and interactions of constituent components in a variety of musical structures, we can represent structure along a shared dimension that reflects the relative transparency of internal relationships. We could define the poles of this dimension with *Concret PH* (unclear) and *My Reverie* (clear) (or other similarly appropriate examples). The poles of this dimension would be open (given they are unbounded) with a neither / nor intermediary. While we don't have the experimental data to more precisely quantify the relative size of the neither / nor region, intuitively, we imagine its relative size is significant, so we can represent the dimension something like the following (Figure 7.8):<sup>57</sup>



**Figure 7.8. Dimension representing transparency of structural relationships.**

*In for steps that grow when climbed*, significant contrast occurs along this dimension contributing to our sense of discontinuity in the music. For example, two entities that vary significantly in their structural clarity are juxtaposed at 3:32.75. The material before this moment is dense, timbrally diverse and registrally expansive. While some components of this texture are discernible, such as the short looped piano sample, the granulated oboe and the noise / bit degraded elements, the texture more than anything projects a particular kind of “wash”. The result is a complex fabric whose identity as a


<sup>57</sup> Dimensional representations throughout this section will be similarly hypothetical.

whole occludes the identities of and relationships between its components. Conversely in the section following 3:32.75, we are presented with a texture whose structural components are significantly more clear. Here, the bassoon plays melodic passages, the tom toms articulate meter and points of structural importance and the piano provides supplementary harmonic material in a dialog with the other instruments. We can represent this discontinuity relation as follows (Figure 7.9):



**Figure 7.9.** Representation of 3:32.75 along a structural transparency dimension. The blue dot represents *entity A*; the orange dot represents *entity B*.

If we were to track this dimension over the course of the work, we get a sense of how structural juxtapositions correspond to moments of overall discontinuity (see Figure 7.10).



**Figure 7.10.** From top to bottom: timeline, perceived discontinuity, acoustic / electronic dimension and structural clarity dimension in *for steps that grow when climbed*.

While many of the structural juxtapositions happen to correspond to changes in the acoustic-electronic dimension, such as at 1:39 and 3:32.75, this is not always the case.

At 4:45, a change in structural clarity corresponds to perceived discontinuity despite a lack of change in the acoustic-electronic vector (this is not to say that this is the only parameter changing at the moment, but it does speak to the important influence of the structural clarity dimension). In the section preceding 4:45, the flute, clarinet, bassoon, viola and double bass are instructed to play as rapidly and as softly as possible, either high or low in their register. The only parameter that changes over this section is phrase length, which starts at seven and incrementally decreases reaching two at the end of the section. Several takes of this section were recorded and then superimposed in the mix. The result is a dense, fairly chaotic texture where the relationships between instruments are not clear (although the timbral and registral identities allow us to decode some relationships, thus, this point is located between unclear / clear). Conversely, in the section after 4:45, the instruments of the ensemble assume specific roles relative to the themes, rhythms and interactions of the section, conveying a structure that is more clear. The resulting contrast is dramatic, and is a key factor in the discontinuity perceived at this moment.

#### *Beat-based – non-beat-based*

The final primary parameter of the piece that I will identify is characterized by the presence of a beat. There are sections that are strongly beat-based, sections where beat is essentially indiscernible and sections in between these extremes. As these sections are juxtaposed with each other, varying degrees of (dis)continuity are produced. Let us first consider how we can think of and represent beat presence as an attribute.


We can describe the presence of a beat as a parameter that can be represented on a dimension. Now, one may be tempted to describe the presence of a beat as a feature, given it is either there or it is not. Upon further reflection though, it satisfies the requirements of a dimension, as described in chapter I, for three reasons. First, there can be intermediary values of beat salience, such as when beat-based material is sampled

and manipulated, as in the transition section starting at 0:38.7. Second, these values of beat salience are mutually exclusive. Third, even if the value of beat salience is zero, the potential of a beat is still apparent. Thus, we can describe beat salience in terms of a parameter, which we can represent on a dimension whose poles are defined by non-beat-based and beat-based. We can define beat based and non-beat-based as the maximum and minimum of the scale: if the music conveys a clear beat, that beat cannot be made any more clear. If the music does not have a beat, it cannot be made to have less of a beat. We can then describe both of these poles as half-closed intervals with an open interval as the intermediary state. The dimension could be represented as in Figure 7.11:



**Figure 7.11.** Beat salience dimension in *for steps that grow when climbed*.

The perception of this dimension over the course of the piece is represented in Figure 7.12:



**Figure 7.12. From top to bottom: timeline, perceived discontinuity, acoustic / electronic dimension, structurally clear / structurally unclear dimension, beat-based dimension in *for steps that grow when climbed*.**

Perception along this dimension largely inhabits the poles over the course of the piece, which is not totally surprising given we generally gravitate towards binary judgments of beat presence. There are moments that inspire intermediary interpretations such as 0:38 to 0:43 and 2:10 to 3:10. In both of these moments, the entrance into this intermediary territory is enabled by sampling and processing of earlier beat-based material. As previous moments are elongated and superimposed, structural characteristics, such as a beat grid, that were previously apparent are obscured (this is a primary means of transformation in the piece).

Abrupt changes from beat-based to non-beat-based or vice-versa are associated with significant discontinuities, such as at 1:39, 3:32 and 4:07. While we cannot single out changes in these dimensions as the cause of discontinuity percepts given they occur simultaneously with changes in the previously mentioned dimensions, the juxtaposition of beat-based and non-beat based passages is one of the more striking aspects of the piece, so it is safe to assume that changes along this dimension contribute meaningfully



to the moments of discontinuity noted.

### *Form*

Discontinuity is used to demarcate form and explore relationships between disparate materials. As the piece develops, the discontinuities become more sparse, thus elongating the work's sections. This sort of large-scale configuration experiments not only with how we perceive and remember formal sections of different lengths, but also with how contrasting materials affect such percepts. In terms of the relationships between the work's components, the work explores how different worlds can co-exist and interact to create new textures. As a result, parts are related to those that sounded previously, even though significant discontinuities may separate them. We then recognize connections between temporally separate periods, such as between the highly structured instrumental sections at 0:31 and 3:34, and we understand how disparate materials can interact with each other, such as when the instrumental section at 4:10 adopts the rich, chaotic density of the electronic texture of the section that occurs before 3:32. In regards to the configuration of these various kinds of entities, the piece gets from one place to another in a variety of ways and thus experiments with how different kinds of transitions affect (dis)continuity perception. For example, at 3:34 two contrasting entities are juxtaposed with a silent intermediary, while from 2:10 to 3:07, a longer, linear transition makes a connection between materials that in isolation, are quite unlike. Finally, the work shows how we privilege particular high-level parameters, acoustic / electronic, structured / free and beat-based / non-beat based in this piece, and how changes along those parameters can significantly influence our interpretation of the piece's form.

## Breeding in Pieces

2 channel recording, 2009

Contrast can exist in both continuous and discontinuous musical passages: *Breeding in Pieces* explores the possibilities of this idea. One of the factors that plays a role in whether contrast will be perceived as continuous or not involves the temporal characteristics and configuration of entities. The work exhibits how the juxtaposition of brief but significantly contrasting entities can be perceived as continuous through the idea of *micro-discontinuity*. *Breeding in Pieces* also uses contrast to create perceptual discontinuities, which I will describe by looking at a few salient discontinuous moments, including the most significant discontinuity that occurs at 3:14.

### *Contrast that doesn't create discontinuity*

*Breeding in Pieces* experiments with the idea that contrasting elements that contain the right characteristics, are configured in a particular way, and function in an appropriate context will not produce structural segmentation boundaries or a sense of discontinuity. The beginning of the piece exhibits this postulation (Figure 7.13):

[Click to listen](#)

**Figure 7.13.** *Breeding in Pieces* 0:00-0:07.

From a compositional perspective, the five chords that are contained in this excerpt differ in terms of pitch configuration, instrumental source, number of voices, spectral character, stereo position and processing. For example, a bountiful amount of delay is applied to the first chord while others are completely dry. In the language of structure mapping theory, the opening of the piece is a string of analogies. Functional relations are shared between the various chords, but their attributes differ significantly.

Interestingly, we hear the sequence of chords as continuous despite these differences. This supports the aforementioned argument that we privilege 1) relations over attributes and 2) higher-level constructs, such as “a guitar playing a chord progression” over their lower-level counterparts, such as stereo positioning or timbral character. These strong relational connections engender a sense of connectedness and thus, continuity. The following section, 0:10-0:20, functions to clarify the heterogeneity *yet* continuity of the opening section by featuring a guitar progression with few attribute changes (other than oscillating reverberation).

The section from 0:49 to 1:22 also exhibits contrasts that don’t produce judgments of discontinuity (Figure 7.14).

[Click to listen](#)

**Figure 7.14. *Breeding in Pieces* 0:49 to 1:22.**

This excerpt explores the idea of simultaneous discontinuity mentioned in chapter III. The guitar parts provide a continuous flow of sonic information over which three significantly differing drum parts are superimposed. The drums, as a whole, convey discontinuity in two primary ways. First, each drum part is separated by significant silence: 1.8 seconds separates drum parts one and two and 5.5 seconds separates drum parts two and three. Second, each drum part differs in terms of timbre, stylistic association, dynamics and rhythmic configuration. Drum part number one is dynamically soft and bandpass filtered, creating a “telephone” effect. Drum part number two appears in full fidelity but is still mixed relatively low. It also uses all electronic sounds, evoking an association with electronica drum beats. Drum part number three is also in full fidelity, but is mixed much louder (as compared with number one and two), consists of fully acoustic drum sounds and evokes stylistic

association with rock. The rhythmic configuration differs between each drum part, but those differences are less salient than the ones previously mentioned.

Context is one of the main reasons that this passage is not perceived as discontinuous. Considering the sequence of these drum parts out of context, there are plenty of ingredients sufficient to create discontinuity between these entities, as described above. However, in context, these parts don't appear as particularly discontinuous. One reason, as in the previous example, is that these drum parts are analogies: they differ in terms of several attributes, but their relations relative to the rest of the music are almost identical (for example, the drum parts in all cases articulate the rhythmic and metric structure of the guitar parts. Additionally, the relationships between the constituents of each drum part are highly similar, such as between the tom and the snare drum). Because of the relational similarity, almost all of the differences are alignable. The strong relational similarities exert primacy over the contrast in attributes, inspiring a sense of connectedness and continuity. The other reason is that the guitar part, which plays a meaningful role in the experience of this section, is generally continuous. The combination of these two piece-level attributes overwhelm the lower-level discontinuities mentioned to inspire a general sense of continuity when listening to this section.

### *Micro-discontinuity*

The section from 1:24 to 1:53 exhibits a different kind of configuration that re-orders an entity into temporally abbreviated, contrasting elements yet maintains a sense of continuity; I will refer to such configurations as *micro-discontinuity* (Figure 7.15):

[Click to listen](#)

**Figure 7.15. *Breeding in Pieces* 1:24 to 1:53.**

The section begins with the statement of a guitar theme. As the theme repeats, tiny segments are incrementally removed, added and re-configured in a way that is analogous to cubist painting processes. With each repeated iteration of the theme, more and more segments are reconfigured, so that by the end of the section, a significant portion of the music is constituted by these reconfigured segments.

Typically, the type of dramatic contrast created from reconfiguring an entity would inspire notions of discontinuity, but here, that is not the case. Instead, we maintain a sense of the theme despite its cubist reconfigurations, inspiring a sense of continuity. There are two main reasons for this.

First, and most importantly, we heard the theme in an unaltered state at the beginning of the section thus establishing it as a memorable anchor, which enables us to understand how it is transformed in subsequent iterations. The contextual surroundings of the beginning of the section enhance the theme's memorability thus partially enabling this continuity. The previous section, which includes many instruments, comes to a climax and then abruptly halts after which the guitar theme is immediately presented. This moment is perceived as discontinuous and a strong sectional boundary, a function of the contrast in dynamics, instrumentation and number of voices. As a result of this salient discontinuity, our attention is heightened at the moment when the theme is introduced, contributing to its perceptual memorability. The character of the unaltered theme thus becomes knowledge, which is then recalled every time a version of the theme, no matter how transformed, sounds. Because we can recognize the theme in each iteration, we can recognize how that theme changes, which allows us to perceive the transformation that manipulates the theme. The result is that we perceive a sequence that is characterized by a higher-level, abstract *transformational equation* that describes each iteration and thus creates a sense of connectedness and continuity. Because we privilege higher hierarchical levels, the music projects continuity despite its surface features changing significantly.

The second reason that the segmentation and re-configuration of the theme is perceived as continuous involves the temporal character of the segments, which are often smaller than 100 msec and are spaced tightly. Because of such temporal identities, we aren't able to gain a great sense of the character of individual events. That is, we couldn't say from where a particular segment was extracted, or even if a segment was part of the original sound file. Instead, we have a more abstract understanding of these micro-blocks: we know, in a general way, that they are temporally abbreviated and timbrally-varied. It could be the case that because we aren't able to individually differentiate these small segments, we group them together into a higher-level organization that comprises temporally-abbreviated and timbrally-varied segments. We then realize that these higher-level organizations are substituted for parts of the theme, which we are able to do because we have established knowledge of the unaltered version of the theme. Because of this hierarchically-influenced perception, we connect the various iterations and perceive the section as continuous.

*Contrast that creates discontinuity*

The work also contains dramatic contrasts that create strong sectional boundaries and discontinuities. In order to explore what kinds of changes are sufficient to create a sense of segmentation and discontinuity in this context, I will explore the most dramatic example of discontinuity in the piece, which occurs at 3:14 (Figure 7.16):

[Click to listen](#)


**Figure 7.16. Discontinuity at 3:14 in *Breeding in Pieces*.**

The discontinuity here is created by differences in attributes and relations, thus, we can describe this moment as an anomaly in the language of structure mapping theory.

Attribute contrasts, which include dynamics, number of voices, instrumentation,

stylistic association, rhythm and pitch content, are a main reason that we perceive sectional boundaries and discontinuity at 3:14. The dynamic contrast is significant, the RMS of *entity A* (which ends at 3:13) is -18.4 dB while the RMS of *entity B* (which begins at 3:14) is -29.4 dB, or A is 2.14x as loud as B<sup>58</sup>. The number of voices and instrumentation also differ significantly. In terms of number of voices, the two entities are not comparable: A includes three while the number of voices that constitute B are not able to be specified (we only have a sense that B comprises a fairly large number of small, brief elements). In regard to instrumentation, *entity A* includes three main voices: two guitars panned left and right and a drum set, while *entity B* presents an intricate but amorphous electronic texture. Thus, we can describe both of these contrasts as nonalignable differences that resist representation on a shared dimension.

The contrast in stylistic association is more clear, *entity A* evokes rock while *entity B* evokes electroacoustic music, but it raises representational issues. We have difficulty locating such contrast on a single shared dimension, for what poles would define such a dimension? Instead, stylistic differences would have to be measured in a multidimensional space, such as the kind described in Novello et al. (2011) (Figure 7.17):



<sup>58</sup> Loudness values calculated with the tools on [sengpielaudio.com](http://sengpielaudio.com)

**Figure 7.17. Novello et al. (2011) Fig. 7: “Genre topology: two-dimensional MDS plot, calculated from the 13 X 13 dissimilarity matrix constructed from the relative distances of the 13 genre centroids (Pearson correlation coefficient . 0.86). The arrows drawn represent the separation between prevalently acoustic and synthetic genres. The values on the axes refer to the point distances in the two-dimensional MDS-space.” (12).**

I am reluctant to use Novello et al. (2011)’s multi-dimensional space to calculate a value for the discontinuity at 3:14 in this piece because the former was created from popular music excerpts that don’t include the electroacoustic genre. Nonetheless, Novello et al. (2011) gives us an idea about how we can make such quantitative comparisons given appropriate experimental data.

The rhythmic contrast of this moment is also significant. *Entity A* comprises two rhythmic streams that project independent tempi, meters, lengths, and rhythmic configurations (thus the ideas of polymeter, poly-tempi and polyrhythm are appropriate descriptors here). Further, each of the aforementioned rhythmic parameters change regularly within each voice, creating even greater rhythmic diversity. Orchestrationally, one stream is expressed by one guitar, the second by the other guitar, and different parts of the drum set are mapped between the two guitar parts (creating perceptual grouping competition that could be the topic of another discourse). Despite the rhythmic complexity of this section, our sense of some kind of beat is palpable, however may mercurial it may be. In contrast, *entity B* projects none of the rhythmic attributes of *entity A*. *Entity B* does not express beat, meter, tempo, rhythmic configuration or even phrase or section lengths. Thus, as in the case of number of voices and instrumentation, the tempo, metric, length and rhythmic configuration differences between A and B are nonalignable. We can, as we did in *for steps that grow when climbed*, describe beat salience in relation to a parameter, which would be represented as in Figure 7.18:





**Figure 7.18. Representation of beat-based versus non-beat-based contrast at 3:14 in *Breeding in Pieces*. The orange circle represents *entity B* and the blue circle represents *entity A*.**

*Entity B* is represented at the pole because *entity B* is as non-beat-based as it gets. *Entity A* is close to the beat-based pole, but because of the rhythmic complexity and change mentioned earlier, the sense of beat is a bit more ambiguous than those examples that would be characterized by the pole of this dimension. In any case, the result is significant contrast along this dimension, which contributes to our sense of discontinuity at 3:14.

*Entity A* and *entity B* also differ in terms of pitch content. *Entity A* comprises the E G A B C D pitch set (e minor and its transpositions characterize much of the pitched material in the piece). In contrast, *entity B* doesn't project definitive pitches at all, instead appearing as a mosaic of electronic noise elements that are inharmonically related. Because these differences are nonalignable, representing the contrast on a shared dimension is difficult in the language of pitch. Instead, we can compare the spectral content of the two entities to gain some sense of their frequency-related contrast. Figure 7.19 and Figure 7.20 compare the loudness-equalized spectra (*entity B* is amplified so that it is -19.9 dB RMS and -1.3 dB peak vs. *entity A* which is -20 dB RMS and -1.7 peak) of the two entities on a linear and log scale, respectively:




Figure 7.19. Spectrogram of the discontinuity at 3:14 in *Breeding in Pieces*. Y-axis scale displacement: log; window size 4096, overlap 87.5%; frequency bins: linear scale.




Figure 7.20. Spectrogram of the discontinuity at 3:14 in *Breeding in Pieces*. Y-axis scale displacement: linear; window size 1024, overlap 50%; frequency bins: log scale.

From these figures, we can identify a number of defining features for each entity. *Entity A* has regular sonic spikes that cover the frequency range fairly uniformly (Figure 7.19), which are produced by the percussion, and also has strong, harmonically related components from 100 Hz - 2 kHz (Figure 7.20), which are produced by the guitars. Conversely, the frequency distribution of *entity B* is more discrete and is distributed throughout the frequency range in a less harmonic way. These representations help describe the anomalous but significant frequency / pitch-related contrast that contributes to the discontinuity perceived at 3:14.

The relations between entities involved in this moment also play an important role in creating discontinuity. *Entity A* is characterized by the relationship between the two guitar parts and the relationship between the drums and each guitar part. In one sense, the guitars function independently of each other, a judgment that is inspired by their rhythmic characteristics (as previously described) and their wide stereo spread. Because of this independence, each guitar part is considered as important as the other in terms of defining the musical content. The drums, while they may project an identity as an independent whole, in actuality, the components of the drums are mapped to guitar parts in specific ways. That is, the kick drum and hi-hat typically demarcate the phrases of one guitar part and the ride and the snare demarcate the phrases of the other. In this way, the relationship between the percussion and the guitars is quite clear and characterized by simultaneity. At the same time, these relationships are ambiguous because we can hear the drum set as an independent whole, or we can hear the drum set's components align with the guitar parts. This relational ambiguity is one of main ideas of this section, and is one of the reasons that the section is elongated in order to give the listener time to explore these perceptual possibilities. *Entity B*, in contrast, does not contain any of these relations. More simply, we can either perceive *entity B* as a single, complex sound, or we can consider it as a collection of many small,

indistinguishable components that create a complex texture. In either case, the differences between the two are nonalignable.

Bringing these various elements together, the discontinuity at 3:14 is primarily a function of alignable differences and hierarchical position (higher levels are more influential). When I hear the discontinuity at 3:14, I am primarily struck by stylistic contrast, and secondarily by volume and beat contrasts. Indeed, these alignable differences (dynamics, style and beat) seem to be more salient than the more numerous nonalignable differences (number of voices, instrumentation, rhythm and pitch content). This supports the aforementioned claims about how we privilege alignable differences and higher hierarchical levels. In support of the latter, we can consider style as higher-level than either volume or beat given style references a richer network of interconnected knowledge. In sum, the discontinuity at 3:14 is defined as anomalous but with particularly salient alignable differences in style, dynamics and beat, which garner perceptual attention.

### *Hierarchies and Context*

*Breeding in Pieces* explores how (dis)continuity can be perceived on multiple hierarchical levels. The piece uses low-level contrast to create musical textures that are appear continuous on higher levels, such as at 1:24. The piece also uses contrasting materials to create discontinuities that are perceived on all hierarchical levels, which create sectional segmentations. These segmentations often involve stylistic discontinuities (which mostly consist of rock and electroacoustic elements), which segment the work into its formal sections, such as at 3:14. Such segmentations support the notion that higher-level percepts, style in this case, play an important role in how we prioritize perceptual organizations in a listening experience.

The piece also explores how juxtaposing stylistically contrasting entities can shape our understanding of those entities, thus exhibiting how context affects our

musical experience. There are some moments that are wholly rock; there are some that are wholly electroacoustic; and there are moments that are some synthesis of the two. The work recognizes that we remember this diversity as a new section is presented, and that as a result, we may be reluctant to accept any particular gesture as fully authentic. At the same time, the piece tries to present authentic gestures, and invites listeners to consider them as such. Whether they will or not depends not only on the individual qualities of listeners, but also on characteristics of the piece, such as sectional durations, type of content, and the configuration of contrasting entities. More simply, musical passages assume new identities by being placed in a variety of contexts.

### **Push for Position**

For bassoon, saxophone, the robots AMI and CARI, and electronics


2011

I will describe how *Push for Position* approaches discontinuity from the *poeitic* perspective. Because the piece was composed later than both *for steps that grow when climbed* and *Breeding in Pieces*, much of the framework described in the previous chapters had been developed before the piece was written. Therefore, some of the ideas of the framework were implemented during the compositional process. Using a few examples from the piece, I will describe how multi-parametric change was designed to effect perceptual discontinuity. Here, the relevant parameters include machine / human, instrument group (acoustic instruments, robots, electronic sounds), sparse / dense, durationally brief / long, energy, pitch set, register, free / structured, beat, order / chaos, nonsense / sense, timbre, instrumental techniques and stylistic association. From there, I will discuss some of the larger governing ideas of the piece to illuminate how these smaller discontinuities contribute to the piece's overall form. My descriptions will

reference timings as indicated in the score for the piece.

The first discontinuous moment occurs at 0:45. The music prior to this moment, *entity A*, is generally sparse, subdued, chromatic, is articulated in the lower parts of each instrument's register and conveys neither a beat nor rhythmic configurations. At 0:45, the performers are given the instruction "low register freak out" (resulting in *entity B*). An explanation of this instruction: at first, I considered telling the performers which parameters along which to create contrast. I imagined that such an instruction could be overly taxing, so I instead gave them a higher-level instruction that implies change along a number of lower-level parameters, including dynamics (soft to loud), order to chaos, sparse to dense and subdued to energetic (these are the parameters that dictate the contrast in the electronic and robotic parts). Part of the interest in such an instruction is the interpretive flexibility that it gives performers, which can result in the articulation of diverse parameters. In any case, if we assume that the performers create contrast along the aforementioned parameters, *entity A* is located very close to one pole and *entity B* is located very close to the opposing pole within each corresponding dimension. The fact that such high contrast occurs along at least four parameters ensures that a sectional boundary will be created at 0:43. As the music of the human performers, robots and electronics change similarly, the above distinctions are strengthened. The sections at 0:58 and 1:07 function similarly, with the exception that at 1:07, there is also a registral change from low to high. The purpose of this addition is to exhibit how change along one additional parameter, register in this case, can affect the strength and kind of discontinuity relative to the ones at 0:45 and 0:58.

The section starting at 1:14 creates discontinuity along the parameters of machine / human, sparse / dense, free / structured, beat, energy, and pitch set. *Entity A*, the section before 1:14, is articulated by all of the performers, *entity B*, the section from 1:14 to 1:24, is articulated by only the robotic and electronic elements. We could represent this contrast dimensionally, as in Figure 7.21:



**Figure 7.21. Discontinuity at 1:14 represented on a machine / human dimension. The orange circle is *entity B*, the blue circle is *entity A*.**

The poles are bounded points, separated by an intermediary neither / nor open interval. In the section before 1:14, *entity A* is a fairly balanced mix of human and machine elements, but there is a slight weighting towards the machine pole. *Entity B* is only machine elements, thus, it is located at the pole. The contrast is striking but as Figure 7.21 shows, more extreme contrast is possible.

The temporal contrasts between the A and B can be described in terms of sparse to dense, free to structured and non-beat-based to beat based. The sparse / dense dimension can be represented as two open interval poles separated by an open neither / nor interval. *Entity A* is fairly sparse, while *entity B* is quite dense, as seen in Figure 7.22:



**Figure 7.22. Discontinuity at 1:14 represented on a sparse / dense dimension. The blue circle is *entity A*, the orange circle is *entity B*.**

The free / structured parameter functions on a higher hierarchical level as it involves beat salience, pitch set, composed versus improvised and synchronization. We don't need to necessarily represent each of these lower level parameters (though we could) to define this higher-level parameter. Instead, we can simply ask ourselves, or participants, to what extent do we perceive structure in these entities. We can thus represent a free / structure dimension as two open interval poles with a neither / nor intermediate. The

contrast along this dimension at 1:14 here is extreme, as *entity A* is perceived as significantly free and *entity B* is considered significantly structured (see Figure 7.23):



**Figure 7.23. Discontinuity at 1:14 represented on a free / structured dimension. The blue circle is *entity A*, the orange circle is *entity B*.**

The final temporal component is beat-based / non-beat-based, which is admittedly related to the free / structured dimension. We can represent the beat-based / non-beat-based dimension as before, as two half-closed poles and a neither / nor intermediary. As with the free / structured dimension, here, the contrast between *entity A* and *entity B* is extreme, contributing to the overall sense of discontinuity (see Figure 7.24):




**Figure 7.24. Discontinuity at 1:14 represented on a beat-based / non-beat-based dimension. The blue circle is *entity A*, the orange circle is *entity B*.**

Energy, as in the case of the free / structured parameters, can be considered higher-level given it comprises lower-level parameters such as dynamics, rate, and articulation. As in the case with the free / structured parameter, we perceive energy as a whole, and thus, we can measure it along a dimension without specifying its constituents. The section preceding 1:14 is generally low-energy while the section after 1:14 is high energy. If we define the poles of the corresponding dimension in terms of energy level as typically experienced in music, both poles would be characterized by open intervals with an intermediary open neither / nor interval. The representation of



the energy contrast at 1:14 could be represented as in Figure 7.25:




**Figure 7.25. Discontinuity at 1:14 represented on a sparse / dense dimension. The blue circle is *entity A*, the orange circle is *entity B*.**

Finally, pitch contrast also contributes to the discontinuity at 1:14. The music prior to 1:14 is primarily chromatic while the gestures in 1:14 to 1:24 are based on C# Major. In this particular situation, it is very difficult, if not impossible to quantify the perceptual difference between pitch sets because in *entity A*, the instrumental parts are improvised and the pitches are randomly determined.<sup>59</sup> Therefore, it is possible that the performers, consciously or unconsciously, gravitate towards pitch collections that are a subset of the larger chromatic scale. Similarly, it may be, through random pitch determination, that some smaller pitch subsets are perceptually emphasized. Thus, when *entity B* begins, which clearly emphasizes C# Major, the contrast will vary depending what subsets were emphasized before 1:14. While we cannot represent the contrast here (we could for a realization of the piece, but not for the set of perceptual possibilities afforded by the score), we can still say that the contrast between a general sense of a chromatic pitch space and C# Major is significant, contributing to the discontinuity at 1:14.

Combining the above parameters, the aggregate contrast is significant, so it is no surprise that a perceptual segmentation boundary is created at 1:14. While the discontinuity at this moment is indeed powerful, the entities could be more distant within each parameter, so there is the possibility that a more salient discontinuity could be created using these same parameters.

<sup>59</sup> The latter may contribute to our sense of *entity A* as being “unstructured”, but such a claim would need to be validated by more thorough consideration on the connections between randomness and the perception of structure.

A section similar to the one heard in 1:14 to 1:24 returns at 3:28, except this time, it is preceded by a “freak out” section rather than the sparse chromatic improvisatory section heard prior to 1:14. The purpose of this repetition is to explore how the same material can evoke different notions of discontinuity depending on the context. I will consider this moment along the same parameters described in the above paragraphs to highlight these differences. As the parameters and associated dimensions have already been explained, I will provide all the representations in one place (see Figure 7.26):



**Figure 7.26. Discontinuity at 3:24 in *Push for Position* represented along several dimensions. The blue circle is *entity A*, the orange circle is *entity B*.**

As we can see, there is significant contrast along some parameters, such as beat-based / non-beat-based, and free / structured, but there is less contrast in other dimension such as low energy / high energy and sparse / dense compared to the discontinuity at 1:14. Also, note how context affects the perception of these parameters. For example, in the discontinuity at 1:14, *entity B* is perceived as near the *dense* pole, while at 3:24, *entity B* is perceived relatively further away from the same dense pole. The reason for this is the context: *entity B*, when following a sparse *entity A* that precedes 1:14, seems near

maximum density, but *entity B*, when following a more dense *entity A* that precedes 3:24, suddenly doesn't seem quite as near the pole anymore. As a result, not only is the discontinuity at 3:24 not as striking, but also our perception of *entity B* changes relative to its statement in 1:14 as a result of the particular juxtaposition heard in 3:24. Thus, the context of a moment significantly affects how we perceive and represent entities, attributes, relations and thus, contrast.

There are several other stark juxtapositions that are designed to inspire a sense of discontinuity. For instance, at 1:24, a discontinuity is created involving the parameters of human / machine, free / structured, beat-based / non-beat-based and sparse / dense. At 1:44, a discontinuity is created involving the parameters durationally short / long, dynamics and energy. At 4:14, a discontinuity is created involving the parameters pitch set, energy and density. At 6:31, a discontinuity is created involving the parameters of free / structured, rhythmic configuration and stylistic association. At 7:39 a discontinuity is created along the parameters of sense / nonsense, free / structured, dynamics, beat-based / non-beat based and pitch set. At 8:04, a discontinuity is created along the parameters of register, density, dynamics, order / chaos, timbre, instrumental techniques and electronic instrumentation. Each of these moments could be described and represented as in the previous two examples. The compositional intent, as described in the examples above, is to explore how parametric change in terms of type, degree and number, affect discontinuity perception.

### *Form of Push for Position*

With a general sense of the kinds of discontinuities that occur in the piece, I will explore some of the ideas that connect these disparate sections and articulate the work's form. The formal construction of this piece is not simply a sequence of contrasting blocks (though there is much to be explored in that kind of form). Instead, the work is governed by a number of larger ideas that create relations between these contrasting

blocks.

The first idea is that the piece has memory. Once a section has been stated, it will influence the material and organization of future sections. For example, the 1:14-1:24 section introduces a new pitch space: C# Major (the pitch material is taken from the chromatic section in the prior section). When a new and contrasting section is introduced at 1:24, the feeling of the section is reminiscent of the music that preceded 1:14, except now, it “accepted” the parametric revision offered by the 1:14 -1:24 section (C# Major). In addition, the density of the section at 1:14-1:24 is also remembered, which is applied to the previously sparser music prior to the 1:14-1:24 section. In this way, the current music looks into the past and combines parametric qualities from different sections. The section that begins at 5:14 is another example. As the music remembers the section that begins at 2:47, it settles into strings of isochronous repeated pitches.

Another governing idea of the piece is that multiple trajectories are simultaneously present, although we only get a glimpse of one trajectory at a time. As these contrasting trajectories take turns on the musical surface, the piece explores the possibilities of interrupted continuity beyond the A B A paradigm. Here, there are multiple possibilities for *entity A* and *entity B*, so a particular section could be viewed as the continuous material or the interruption. The trajectories are characterized by the scattered, loosely organized material of the beginning of the piece; the high intensity, dense, chaotic passages (such as at 0:45), and the structured, theme driven sections (2:47, 6:31).<sup>60</sup> These trajectories aren’t wholly independent; they are aware of the developments in the other trajectories and change according to what is presented. As a result, while these core trajectories take abrupt turns to appear on the surface, creating discontinuities, they continue throughout the piece, evolving as subsequent music is presented.

---

<sup>60</sup> Note, the thematic trajectory is not characterized by one theme that develops over the course of the piece, rather, it is characterized by the more abstract idea of structured, composed “theme”, which stands in contrast to the other trajectories. Despite this abstract category of “theme”, there are important parametric connections between the various thematic iterations.

As these trajectories influence each other to a greater extent, they flow towards confluence, which brings them to the surface and creates a transition. The section from 4:41 to 6:07 is such an example. Here, the scattered, chromatic and loosely organized trajectory, first heard in the beginning of the piece, conflates with the thematic trajectory so that the registrally scattered qualities of the former are synthesized with the harmony and beat-based qualities of the latter. Thus, this governing idea of inter-trajectory influence takes over block juxtaposition as the method of formal organization in the piece. When this conflation is fully realized, it disperses into silence, leaving space, which then catalyzes the pure statement of the piece's governing theme. The latter shows that when confluence is fully realized, pure statement of a trajectory is then possible. In addition, it shows that the evolution of trajectories is not linear from start to finish. That is, there is room to create discontinuity along such high-level organizations.

In sum, the piece's form is dictated by discontinuities on various hierarchical levels. On all levels, the piece explores how the number, type, degree of contrasts (which I describe in terms of musical parameters) can affect the perception of discontinuity and structural segmentation boundaries. On higher hierarchical levels, the work's form is governed by the inter-relationship between a number of musical trajectories. As these trajectories contrast, they demarcate form as they assume the musical surface, which expand the possibilities of interrupted continuity as previously described. As these trajectories are aware and have memory of their counterparts, their interactions are expressed in ways other than block juxtaposition. As these trajectories stop and start, they create various kinds of discontinuities on various levels.

### **From Here to There**

For AMI and CARI

2011-2012

*From Here to There* explores the idea of transformational distance between contrasting entities, as explained in chapter IV. The entities in this context are five themes (see Figure 7.27):

**Theme 1**  
 $\text{♩} = 175$   
 CARI *ff*  
 AMI *ff*

**Theme 2**  
 $\text{♩} = 150$   
 CARI 3:2 3:2 3:2 3:2  
 AMI

**Theme 3**  
 $\text{♩} = 150$   
 CARI 3 3 3  
 AMI

**Theme 4**  
 $\text{♩} = 170$   
 CARI  
 AMI

**Theme 5**  
 $\text{♩} = 87.5$   
 CARI *f*  
 AMI *f*

Figure 7.27. Five themes in *From Here to There*.

The piece explores discontinuity and transformational distance by juxtaposing and morphing between these various themes. For example, section three of the piece begins by stating Theme 1, which is immediately followed by Theme 2 in measure 4. By

directly juxtaposing such contrasting themes, discontinuity is created. The piece then “measures” the transformational distance between these themes in mm. 4-15.

The piece *explicitly* measures inter-theme contrast by exhibiting the number of steps it takes to transform one theme into another. As I mentioned in chapter IV, we can interpret the idea of transformational steps in a variety of ways. In this work, I consider the number of *compositional* steps it takes to transform one entity into another. Here, compositional steps are higher-level as they comprise many lower-level steps, and not all compositional steps are the same. These steps are designed to be perceptually relevant and musically meaningful, which are qualities that are not present in all transformational distance measurement methods.

This kind of transformational measurement is exhibited in the beginning of section 3, when theme 2 transforms into theme 1 one step at a time. The following describes the transformation between theme 2 and theme 1 from mm. 4 – 15 [*transformation (instrument, measure)*]:

- Tempo  $q = 150 \rightarrow q = 175$  (all, ma. 5);
- triplet eighths  $\rightarrow$  sixteenth notes (CARI, ma. 6);
- rhythmic configuration of theme 2  $\rightarrow$  rhythmic configuration of theme 1 (all, mm. 7-10);
- pitch set of theme 2 [0 1 2 3 4 6 7] (E, F, F#, G, G#, Bb, B)  $\rightarrow$  pitch set of theme 1 [0 1 2 3 5 7] (D, Eb, E, F, G, A) (all, mm. 11-14);
- registral character of theme 2  $\rightarrow$  registral character of theme 1 (all, ma. 15)

Thus, the transformation from theme 2 to theme 1 can be done in five steps, which can be described along the parameters of tempo, rate, rhythmic configuration, pitch set, and registral character. As expressed above, each of these transformational steps is higher-level in that it comprises many lower level steps. I chose to define transformation in this way because explicit lower-level steps, such as changing individual pitches or rhythmic



values, may not result in perceptible change from step to step. Further, such transformations would take many repetitions, which would likely induce habituation and significantly affect the aesthetic character of the music.

### *Describing Transformations*

The process described above produces discernible change at each transformational step; we can describe the relations that constitute such change in terms of physical characteristics and perceptual experience. As both kinds of descriptions can illuminate the nature of discontinuous relations and our experience of them, I will use both of them in the following sections.

In regard to the transformation at the beginning of section 3, we can focus our physical descriptions on rate, rhythmic configuration, and pitch content. In regard to rate, in step 2, we can describe the transformation in terms of specific interval values, which here are 114 msec (triplet eighth at  $q = 175$ ) versus 85.7 msec (sixteenth note at  $q = 175$ ). We can then represent these values on an appropriate dimension and measure the distance between them to illuminate relative (dis)similarity, which we can then compare to perceptual ratings of such changes.

In step 3, the rhythmic configuration transformation, the pitch content and registral contour of the theme 2 is combined with the rhythmic configuration of theme 1. This combination affects a number of lower-level parameters. First, since theme 1 is four measures long and theme 2 is one measure long, the transformational step duration is extended from one measure to three measures. Because all rhythmic values were transformed to align with the *beat / multiple of 2* grid in the previous step, the rhythmic reconfiguration in this step involves adding and deleting rhythmic values in various places. We could then use an edit distance metric to measure the relative distance between the two steps, as described in Orpen and Huron (1992) (see Figure 7.28):

CARI, step 2	1	1	1	1	1	1	1	1	1	1	1	1		
CARI, ma. 1 of step 3	1	1	1	1	1	1					1	1	1	1

**Figure 7.28. Edit distance between CARI part in transformational step 2 and first measure of CARI part in transformational step 3.**

Thus, the difference between rhythmic configurations in Figure 7.28 can be described as four subtractions and two additions. The other two measures of step 3 contain fourteen additions, bringing the total rhythmic configuration edit distance between the CARI parts in step 2 and 3 to twenty. If we weight subtractions and additions as one (Orpen and Huron, 1992), then we can describe the (dis)similarity between the CARI parts in step 2 and step 3 as twenty, which we can then represent on an appropriate dimension. We could then perform the same steps for the AMI part, and the two parts in combination. Alternatively, we could use another means of rhythmic comparison, as suggested in chapter VI. Whatever method we use, we would likely want to account for the durational distance between entities in the two steps, perhaps on a separate dimension.

We can describe step 4, the pitch set transformation, in terms of commonality. The pitch set of theme 2 (E, F, F#, G, G#, Bb, B) has three pitch classes in common with the pitch set of theme 1 (D, Eb, E, F, G, A): E, F, and G. Thus theme 2 has 43% commonality relative to theme 1 while theme 1 has 50% commonality relative to theme 2. If we were to represent pitch set differences according to these proportions, we would likely have to consider the sequence order, which pitch set came first in the piece, to guide our weighting of these values appropriately. Conversely, we could compare the set theory representations of these pitch sets, theme 2 [0 1 2 3 4 6 7]; theme 1 [0 1 2 3 5 7], to measure commonality and derive a measure of (dis)similarity. A third option would be to represent the pitch sets via a pitch class frequency distribution, which would likely

be a good choice here given there are many repeated notes in the actual music. The most appropriate method will be the one that best aligns with perceptual experience.

The preceding lower-level physical descriptions should be used to complement, not replace, higher-level perceptual ones. The former can help us distinguish transformational distance and types of contrast and thus, are helpful. At the same time, we must recognize that high-level descriptions can illuminate the contrast and transformation between entities without further deconstructive explication, and in some situations, may reflect perceptual experience more effectively than lower-level descriptions. For example, when most people hear theme 2 transforming into theme 1, it is likely that they will have a general notion of pitch set change, and not that a [0 1 2 3 4 6 7] pitch set transformed into a [0 1 2 3 5 7] pitch set. We thus can give a meaningful description of the transformational distance between theme 2 and theme 1 (note, I am not assuming symmetry, that is, the distance between theme 2 and theme 1 is not necessarily same as the distance between theme 1 and theme 2) as five high-level steps. To further specify the degree of transformations, we could ask participants to rate contrast within each of these parameters. We could then compare such responses to the physical analyses previously mentioned to illuminate the relationship between physical and perceptual interpretations. Indeed, the most comprehensive approach will include both physically- and perceptually-oriented methods. Using both will allow us to ensure that our microscopic efforts are conveyed in a perceptually meaningful way.

### *Theme 3 → Theme 5*

Another example of transformation occurs when Theme 3 is transformed into Theme 5 in mm. 74-91 of section 3. As before, I will describe the compositional steps taken to transform theme 3 into theme 1 in terms of musical parameters:

- pitch set (AMI, beat 1 ma. 75)

- tempo (all, ma. 76)
- pitch set (CARI, beat 1 ma. 77)
- pitch set (AMI, beat 2 and 3 ma. 78)
- pitch set (CARI, beat 2 and 3 ma. 79)
- pitch set (AMI and CARI, beat 4 ma. 80)
- rate (AMI, ma. 81)
- rate / configuration (CARI); density (AMI) (ma. 83)
- rate (AMI, ma. 85)
- time signature (all); density / configuration (CARI) (ma. 87)

Relative to the transformation from theme 2 to theme 1 in mm. 4-15 mentioned above, the transformation here takes twice as many steps (10). Part of the reason for this is that transformations are applied to only part of a measure in some cases, which produces more steps. Unlike the transformation from theme 2 to theme 1 in mm. 4-15, multiple transformations can occur in one step here, such as in step 10. The following will explore each step in more detail, considering both physical descriptions and perceptual interpretations.

From a compositional and physical standpoint, many of the transformational steps involve pitch set changes, such as in the first step, which transpires in AMI's part between the first beat of ma. 74 and the first beat of ma. 75. The pitch set of the corresponding part in preceding measure, ma. 74, is [0 2 5 6 8] (E3 G3 B3 C#4 E4 F4 G4), which contrasts of that in ma. 75, which is [0 3] (F#4 F#3 A3 F#4 F#3 F#4 F#3). In regard to the relative similarity of these two pitch collections, we could say that no pitch classes are shared between the two sets, thus, they are significantly contrasting. We could also consider the two pitch sets as complements that together form E F F# G A B C#. Such a combination could inspire associations with D Major or E Dorian, although in this case there is no D pitch class, and F natural would have to be considered a non-harmonic tone. Such an interpretation could still produce contrast given different parts of the larger set are juxtaposed, but reference to a larger pitch structure would likely dilute the

degree of contrast perceived. We could also weight the pitch sets according to frequency of pitch class. In ma. 74, E and G are featured twice as much as B, C# and F. In ma. 75, 85.7% of the pitch classes are F#, so clearly that pitch class defines that part of the measure. In this case, the transformation is primarily between E-G and F#. Again, we could say that no pitch classes are shared between the two, or we could consider them as complements of a chromatic cluster. Alternatively, we could also include A with F#, which is reasonable given A's prominence is amplified by the fact that it is the only non-F# in the segment, which would define the contrast between the two measures as an e minor chord versus an f# minor chord.

With such a variety of possible physical descriptions, we should consider the perceptual experience of the passage to direct us towards the most appropriate interpretation. When listening to the music, the first beat of measure 74 strikes us as a gesture that ascends in pitch more than anything else. Because of the speed of the passage, the individual pitch content is generally obscured, instead presenting a contour that ascends rapidly. While the speed of the corresponding passage in ma. 75 is equivalently fast, the predominance of F#s makes the pitch identity of the entity more clear. In this sense, the contrast involves lack of clarity / clarity more so than a collection of specific pitches. The interpretation that best illuminates our experience of this moment may be some combination of the above descriptions. We could interpret the other pitch set transformations in a similar manner.

The transformation between theme 3 and theme 5 also occurs via rate changes. The first such change occurs in transformational step #7 in AMI's part between ma. 81 and ma. 82. In regard to the rate of notes, ma. 81 largely consists of 32<sup>nd</sup> notes (of the 25 notes, 18 are 32<sup>nd</sup> notes). At a tempo of quarter = 87.5, the duration of a 32<sup>nd</sup> note is 85.7 msec. In ma. 82, the durations of all notes are doubled so that 16<sup>th</sup> notes are the predominant rhythmic value (this also makes the one measure passage a two-measure passage). At a tempo of quarter = 87.5, the duration of a 16<sup>th</sup> note is 171 msec.

Transformational step 9 in AMI's part, where note values are elongated from sixteenth notes to eighth notes, works similarly. Both transformational steps are characterized by slowing down the rate of the passage by doubling the length of notes.

Another rate transformation, step 8, transpires in CARI's part between mm. 81-2 and mm. 83-84. There are two parts to this rate transformation. The first part occurs in the first two beats of ma. 81 and the last two beats of ma. 82, where 16<sup>th</sup> note triplets are transformed into sixteenth notes (114 msec to 171 msec: an elongation). The other part of the transformation occurs between beat three of ma. 81 – beat two of ma. 82 and the corresponding metric position in mm. 83-84. This transformation is not simply a matter of rate change, rather, rhythmic configuration is re-interpreted on a new subdivision grid. The configuration *note-rest-note* is transformed from an eighth-note triplet grid (beats three and four of ma. 81) to a sixteenth note grid (beats three and four of ma. 83), and the configuration *rest-note-note-rest-note-rest* is transformed from an eighth-note triplet grid (beats one and two of ma. 82) to the configuration *rest-note-note-note-rest* on a sixteenth note triplet grid (beats one and two of ma. 83). Thus, this part of the transformation involves changes in both configuration and rate (the rate changes in this part of the passage are accelerations as triplet eights become sixteenth notes). In addition to CARI's rate transformation, the density of AMI's part increases, where sixteenth notes are added to make the rhythm isochronous. In combination, transformational step #8 involves two different rate changes in opposite directions, rhythmic configuration changes and density changes. One would imagine this transformational step would be particularly salient given the number of changes involved.

How these transformations are perceived relative to their physical descriptions presents an interesting question. Halving the rate of notes is a straightforward and familiar operation, but we cannot assume that the degree of perceived change equals that of physical change. Perhaps contextual factors, such as the fact that the tempo

change in step two, quarter = 150 to quarter = 87.5, is not a division by two may influence such percepts. It also may be the case that we perceive rate changes differently depending on the parts of the tempo spectrum that characterize such changes. For example, because the rates in the beginning of this transformation are so fast, our experience consists more of a general sense of speed and less of an apprehension of individual durations. When these fast rates slow into areas where we can apprehend individual durations, our perception of the degree of such change may differ from the physical actuality. The relationship between perceived degree and physical degree may change over the tempo spectrum.

Taking all the transformational steps together, the distance between theme 3 and theme 5 is significant, indicating that the two themes are quite dissimilar. Indeed, if we were to directly juxtapose the two themes, we would likely come to this conclusion. As before, these measurements are produced by high-level compositional choices, and thus could differ in different compositional contexts. This does not invalidate the distance as described, rather, it validates the possibility that different transformational algorithms of different lengths could accurately describe the same transformation. It may be the case that each of these transformational algorithms corresponds to a different perceptual interpretation. Using physical descriptions to complement and explain such perceptual interpretations will enable us to further understand the multi-faced nature of (dis)similarity and (dis)continuity relationships.

### *Form*

As previously mentioned, the piece explores (dis)continuity via juxtaposition and transformational distance between contrasting entities. Other explorations of transformation, such as between theme 2 and theme 4 in mm. 43-57 could be considered similarly as in the above examples. In addition to transformational distance, ideas, particularly theme 1, are expanded, manipulated, and developed according to

compositional intuition and aesthetic considerations. Thus, an artistic context surrounds an experimental exploration. This sort of combination of worlds epitomizes the compositional and theoretical efforts of this discourse.



## 8. Conclusion and Future Questions

### The multi-faceted nature of musical (dis)continuity

Our experience of musical (dis)continuity is complex and is influenced by a number of physical, cognitive and perceptual factors. We can navigate this complexity by considering how we bring distinct entities into relations, for our ability to distinguish and relate is at the core of our ability to perceive musical (dis)continuity.<sup>61</sup> An important way that we can describe how one entity relates to another is by specifying the (dis)similarity of intra- and inter-entity attributes and relational structures (which we can describe in terms of *type* and *degree*). In addition, holistic associations, hierarchical orientation, the content and temporal qualities of intermediaries, and context all play an important role in our (dis)continuity percepts. These various elements work in isolation and in combination depending on the particular musical situation. Importantly, we must recognize that these constructs result from the interplay of physical qualities and perceptual and cognitive interpretations, so it is important that we keep all in mind when describing (dis)continuities. Such recognition elevates the need for some degree of real-time analysis that captures our experience in the context of the music, and not just the context of the score.

### A flexible approach

Because (dis)continuity is multi-faceted, it is unlikely that one of the singular approaches described previously will illuminate the full extent of its richness and complexity. Instead, we must be able to describe (dis)continuity in a number of different

---

<sup>61</sup> At this point, given this topic has been discussed in a number of different ways in the paper, I can now use the singular *ability*, for our ability to perceive likeness and difference, to distinguish and connect, to *relate*, occurs in a way that is synthesized and all at once. This singular quality makes the job of tracing causal connections difficult: do we distinguish by relating? Relate by distinguishing? Do we perhaps do both simultaneously? The possibility of the latter shows the binary constructs that are associated with the products of such judgments, such as *likeness* and *difference*, to be misleading, given they infer that two processes are acting. Such singularity explains why we in many cases have difficulty articulating the interaction of likeness and difference in distinction-making.

ways; the way that we choose depends on the specific musical scenario. In some cases, change along certain musical attributes, such as rhythmic configuration, pitch set or timbre, contributes the most to our perception of (dis)continuity. In others, contrast is created within a particular structure, or by juxtaposing two different types of structures. The tools that I have described, which include holistic associations, structure mapping, transformational distance, hierarchical orientation and dimensional representation, can be used in isolation or in combination to address the facets of (dis)continuity that are most relevant in our musical experience. As experimental evidence shows us how these tools are useful, our models can become more integrative and comprehensive.

A flexible approach is necessary given context can meaningfully affect how we perceive (dis)similarity and (dis)continuity. A particular contrast could be perceived as dissimilar and discontinuous in one musical setting similar and continuous in another. A listener with extensive knowledge of the stylistic and structural conventions of a particular musical style may perceive (dis)similarity and (dis)continuity differently than a listener without such knowledge. This makes the job of constructing models and algorithms that reflect (dis)similarity and (dis)continuity judgments considerably more complex. Instead of being discouraged by such a realization, we can use it to structure our future research of the topic. For example, when we notice differences in (dis)similarity or (dis)continuity judgments, we can look more closely to see what the reasons for such divergences are in regard to differences in listener knowledge and musical context. We can then model different kinds of interpretive scenarios rather than assuming that varied responses reflect that there is no consistency in how people perceive musical (dis)similarity and (dis)continuity.

### **Open Questions**


This paper has detailed a number of questions that could be explored in future research projects. The topics of such questions include auditory and musical contrariety,

the interaction of musical parameters, and our apparent high-level sensitivities. How we perceive auditory and musical contraries, and how we use those contraries to define parametric dimensions is clearly of importance, given the latter has been the primary vehicle that I have chosen to represent relational (dis)similarity. Applying Savardi's (2008) experimental method to diverse musical examples will give us a greater sense of the parameters that function in musical (dis)similarity judgments. As we define the breadth and topography (the latter corresponds to relative salience) of a multi-parametric space, we can explore how various parameters interact in musical relations. This point is of primary importance for multi-parametric interaction is a quality of almost all music. As we understand the relative salience and interactions of these parameters, we can start to understand how the hierarchical level(s) that they, and their associated attributes, inhabit affect our perception of (dis)continuity. As these representations become better defined, their quantifications will become more meaningful. These quantifications can eventually lead to models that allow us to analyze and compose (dis)continuous music more thoughtfully. When asked about the relation between discontinuous entities, we will then be able to say more than "they are unlike."

## Appendix A: Other rhythmic dimensions

### *Durations*

The duration(s) of constituent elements can be quantified using either the actual time value of sounds or some proportional representation of them. Figure A.1 shows how durational discontinuity could be represented via Western notation or as a string of proportions.




**Figure A.1. Western and durational representation of durational discontinuity.**

If the durations within each entity are the same, as they are in Figure A.1, we could represent the durations as either temporal proportions or absolute temporal durations on appropriate dimensions. If the entities comprise a variety of different durations, the situation becomes more complex. One approach to this problem would be to find a representative durational exemplar within each entity and use that exemplar to represent each entity's durational identity as a whole. Alternatively, we could define a dimension beforehand and then ask participants to locate the durational identity of each entity adjacent to a segmentation boundary on that dimension. The distance between these representations would correspond to their relative (dis)similarity and (dis)continuity. The latter would be useful in that it reflects a general sense of a durationally heterogeneous entity without tethering it to any phenomenally present

value or smoothing its inherent variety with a mathematical operation that might not be perceptually relevant.

### *Rhythmic Configuration*

We can represent rhythmic configuration in a number of ways. Toussaint (2004) gives eight such ways of representing the *clave Son* rhythm (Figure A.2):



**Figure A.2. Eight rhythmic representations. Toussaint (2004), p.2.**

We can also represent a rhythm geometrically, via a two-dimensional graph (Gustafson 1987, 1988; Hoffman-Engl, 2002) or a convex polygon (Toussaint, 2004). How we represent a rhythm will affect how we compare it to other rhythms in order to quantify relative (dis)similarity. Because of the many choices involved in both representation and comparison, several models have been proposed for quantifying rhythmic configuration and measuring rhythmic similarity including the Hamming distance, which indicates the number of places where elements in two rhythms differ; the Euclidean interval vector distance, which represents rhythms as vectors of inter-onset intervals that can



If this measure were to repeat many times, thus creating a configuration that consists of the entire measure, how could we talk about rate that characterizes the whole configuration? We could identify the subsections where rate is constant, and then combine those values to produce an overall figure. How would we combine these values? Averaging the rates of the various sections will not give us a sense of the changes of speed that we perceive while listening to the rhythm. Statistical methods such as spike train analysis could better reflect the character of rate change present in a such a rhythm.

### *Syncopation*

Given that rhythms can project various degrees of syncopation, the latter is a dimension along which (dis)continuity can be created. The Harvard Dictionary of Music defines syncopation as “A momentary contradiction of the prevailing meter or pulse” (861). Cooper and Meyer (1960) define syncopation as “a tone which enters where there is no pulse on the primary metric level (the level on which beats are counted and felt) and where the following beat on the primary metric level is either absent (a rest) or suppressed (tied)” (100). Thus, we can understand syncopation as a departure from some foundational structure, such as a beat grid or accent pattern, which, because of a number of psychological factors including rhythmic entrainment and expectation, strikes us as a “hiccup” (this is just one simplified explanation of a complex topic that is best explored in a separate discourse). Several measures have been developed to measure the degree of syncopation present in a rhythm. Gomez et. al (2007) evaluated a number of these measures and concluded that Longuet-Higgins and Lee (1984), Palmer and Krumhansl (1990), Toussaint (2002), Smith and Honing (2006) and Gomez et al. (2005) produced results closest to human judgments. One could represent and compare the output of one of these, or other, methods on an appropriate dimension defined by the characteristics of the applied method.

### *Rhythmic Complexity*

Intuitively, we are able to segment rhythms according to their relative complexity. The problem of defining and quantifying the perceived complexity of musical patterns has been addressed by a number of authors. We can apply the idea of *information-based complexity* (Pressing, 1997), which, as previously stated, is defined as the minimal cost of deriving a solution after identifying relevant variables and describing available operations. In the context of a rhythm, cognitive costs are incurred through maintaining a metrical framework and reconciling the pattern's placement within that framework. The latter corresponds to the handling of syncopation (we could use the syncopation measures earlier mentioned in this context). We can quantify these cognitive costs to derive a measure of complexity.

Shmulevich and Povel (2000) measure the perceptual validity of three approaches that quantify temporal pattern complexity. The first is based on the work of Tanguiane (1993), which describes patterns in terms of smaller constituent patterns and thus measures hierarchical complexity. The second is based on the work of Lempel-Ziv (1976), which highlights the number of “new” substrings as a rhythm unfolds. When a new substring is encountered, the complexity increases by one. As the measure is sensitive to repetitions on all structural levels, it can illuminate dynamic and hierarchical complexities. The third measure, the PS-Measure (developed by the authors), is based on the ideas that a listener attempts to establish an internal clock that corresponds to a rhythmic configuration. Complexity is measured as a function of the induction strength of the best clock and the efficiency of coding the rhythm (Shmulevich and Povel, 2000). By comparing the predictions of these measures with the responses from an experiment that asked participants to judge the relative complexity of rhythmic patterns, the authors found the PS-Measure to be the most robust measure of complexity of temporal patterns. Whatever measure we use, we recognize that a notion of complexity is a function of



many of the earlier mentioned parameters including number of events, rhythmic configuration and beat grid, so we can consider it higher-level. Further, it is likely that the perception of complexity is highly relative, and thus depends on contextual factors previously mentioned, such as surrounding musical textures, stylistic knowledge and musical training. With all of that said, we can create (dis)continuity by juxtaposing simple and complex entities, so we would like to be able to locate entities along a dimension with simple and complex as its poles. We can do so by applying the aforementioned (or other) measures in a way that is sensitive to context and dimensional representation as previously described.

## Works Cited

- Andriessen, Louis. "Zilver." *Zilver*. Perf. The California Ear Unit. New Albion Records, 2009. Audio Recording.
- Applebaum, Mark. *Intellectual Property I*. Program Note. 25 Jan. 2011. E-mail.
- Ballantine, C. "Charles Ives and the Meaning of Quotation in Music." *The Musical Quarterly* 65.2 (1979): 167–184. Print.
- Berenzweig, A. et al. "A Large-scale Evaluation of Acoustic and Subjective Music-similarity Measures." *Computer Music Journal* (2004): 63–76. Print.
- Bianchi, I., and U. Savardi. "The Perception of Contraries." Aracne, 2008. Print.
- Bianchi, Ivana, Ugo Savardi, and Michael Kubovy. "Dimensions and Their Poles: A Metric and Topological Approach to Opposites." *Language and Cognitive Processes* (2011): 1–34. Web. 3 Oct. 2011.
- Birtwistle, Harrison. *Carmen Arcadiae Mechanicae Perpetuum*. London: Universal Edition, c1981. Print.
- . "Carmen Arcadiae Mechanicae Perpetuum." *A/rhythmia*. Perf. Alarm Will Sound. New York, N.Y.: Nonesuch, 2009. Audio Recording.
- Bregman, Albert S. *Auditory Scene Analysis: the Perceptual Organization of Sound*. 2nd MIT Press paperback. Cambridge, Mass: MIT Press, 1999, 2001.
- Brown, T. A. *Afro-Latin Rhythm Dictionary: Handy Guide*. Alfred Pub Co, 1984. Print.
- Cambouropoulos, Emilios. "How Similar Is Similar?" *Musicae Scientiae* 13.1 suppl (2009): 7–24. Web. 3 Oct. 2011.
- . "Melodic Cue Abstraction, Similarity, and Category Formation: A Formal Model." *Music Perception* 18.3 (2001): 347–370. Print.
- Carlsen, J. C. "Some Factors Which Influence Melodic Expectancy." *Psychomusicology: Music, Mind and Brain* 1.1 (2008): 12–29. Print.

- Carter, Elliott. *Sonata for Flute, Oboe, Cello, and Harpsichord* (1952). New York:: Associated Music Publishers, c1962. Print.
- . "Sonata for Flute, Oboe, Cello & Harpsichord." *Elliott Carter: Sonata for Flute, Oboe, Cello & Harpsichord; Sonata for Cello & Piano; Double Concerto for Harpsichord*. Perf. Contemporary Chamber Ensemble. Nonesuch, 1992. Audio Recording.
- Chupchik, G.C., M. Rickert, and J. Mendelson. "Similarity and Preference Judgements of Musical Stimuli." *Scandinavian Journal of Psychology* 23 (1982): 273–282. Print.
- "Continuity." Oxford English Dictionary." Web. 17 Jan. 2011.
- Cooper, Grosvenor. "The Rhythmic Structure of Music." 1963. Print.
- Coyle, E. J., and I. Shmulevich. "A System for Machine Recognition of Music Patterns." *Acoustics, Speech and Signal Processing, 1998. Proceedings of the 1998 IEEE International Conference On*. Vol. 6. IEEE, 1998. 3597–3600 vol. 6. Print.
- Darwin, C. J. "Simultaneous Grouping and Auditory Continuity." *Attention, Perception, & Psychophysics* 67.8 (2005): 1384–1390. Print.
- Deliege, I. "A Perceptual Approach to Contemporary Musical Forms." *Music and the Cognitive Sciences: Proceedings from the Symposium on Music and the Cognitive Sciences, 14-18 March 1988, Centre National D'art Et De culture 'Georges Pompidou', Paris, France*. Routledge, 1989. 213. Print.
- . "Similarity Relations in Listening to Music: How Do They Come into Play?" *MUSICAE SCIENTIAE* 11.1 (2007): 9. Print.
- Deliege, I., and A. El Ahmadi. "Mechanisms of Cue Extraction in Musical Groupings: A Study of Perception on Sequenza VI for Viola Solo by Luciano Berio." *Psychology of music* 18.1 (1990): 18–44. Print.
- Deliege, Irene. "Similarity Perception \_ Categorization \_ Cue Abstraction." *Music Perception* 18.3 (2001): 233–243. Print.
- Deutsch, D. "Auditory Pattern Recognition." *Handbook of Perception and Human Performance*. Vol. 2. 2 vols. New York: John Wiley and Sons, 1986. 32–1. Print.

- Deutsch, Diana. *The Psychology of Music*. New York: Academic Press, 1982. Print.
- "Discontinuity." Dictionary.com. Web. 17 Jan. 2011.
- "Discontinuity." Oxford English Dictionary (oed.com). Web. 17 Jan. 2011.
- "Discontinuity." Answers.com. Web. 17 Jan. 2011.
- Duran, Duran. "Ordinary World." *Greatest*. Capitol, 2004. Audio Recording.
- Garner, W. R. "Aspects of a Stimulus: Features, Dimensions, and Configurations." *Cognition and categorization* 99 (1978): 121. Print.
- Gentner, D. et al. "Similarity Is Like Analogy." *Similarity* (1995): 111–148. Print.
- Gentner, D. "Structure-mapping: A Theoretical Framework for Analogy." *Cognitive science* 7.2 (1983): 155–170. Print.
- Gentner, D., and A. B Markman. "Structure Mapping in Analogy and Similarity." *American Psychologist* 52.1 (1997): 45. Print.
- Girl Talk. "Pump It Up." *Unstoppable*. [Bloomington, Ill.]; [S.l.]: Illegal Art; Spasticated, 2010. Audio Recording.
- Goldstone, R. L. "Similarity, Interactive Activation, and Mapping." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20.1 (1994): 3. Print.
- Gómez, F. et al. "Mathematical Measures of Syncopation." *Proc. BRIDGES: Mathematical Connections in Art, Music and Science*. Citeseer, 2005. 73–84. Print.
- Gómez, F., E. Thul, and G. Toussaint. "An Experimental Comparison of Formal Measures of Rhythmic Syncopation." *Proceedings of the International Computer Music Conference*. 2007. 101–104. Print.
- Goodman, N. "Seven Strictures on Similarity." *Problems and projects* (1972): 437–447. Print.
- Grey, J. M. *An Exploration of Musical Timbre*. Department of Music, Stanford University, 1975. Print.
- Guan, Y., X. Wang, and Q. Wang. "A New Measurement of Systematic Similarity." *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*

- 38.4 (2008): 743–758. Print.
- Gustafson, K. "A New Method for Displaying Speech Rhythm, with Illustrations from Some Nordic Languages." *Nordic Prosody IV* (1987): 105–114. Print.
- . "The Graphical Representation of Rhythm." *PROPH) Progress Reports from Oxford Phonetics* 3 (1988): 6–26. Print.
- Hahn, U., N. Chater, and L. B Richardson. "Similarity as Transformation." *Cognition* 87.1 (2003): 1–32. Print.
- Hasty, C. F. "On the Problem of Succession and Continuity in Twentieth-century Music." *Music Theory Spectrum* 8 (1986): 58–74. Print.
- Helson, H. "The Fundamental Propositions of Gestalt Psychology." *Psychological Review* 40.1 (1933): 13–32. Print.
- Heyduk, R. G. "Rated Preference for Musical Compositions as It Relates to Complexity and Exposure Frequency." *Attention, Perception, & Psychophysics* 17.1 (1975): 84–90. Print.
- Hofmann-Engl, L. "Rhythmic Similarity: A Theoretical and Empirical Approach." *Proceedings of the Seventh International Conference on Music Perception and Cognition*. 2002. Print.
- Huron, David. *Sweet Anticipation: Music and the Psychology of Expectation*. Cambridge, Mass: MIT Press, 2006. Print.
- Ives, Charles. "Central Park in the dark." *Symphony no. 2 The gong on the hook and ladder; Tone roads: no. 1; Hymn; Hallowe'en; Central Park in the dark; The unanswered question*. Cond. Bernstein, Leonard. Perf. The New York Philharmonic. Hamburg: Deutsche Grammophon, 1990. Audio Recording.
- Kessler, E. J, C. Hansen, and R. N Shepard. "Tonal Schemata in the Perception of Music in Bali and in the West." *Music Perception* (1984): 131–165. Print.
- Kim, Chris. "Bio of Alfred Schnittke." <http://www->

- [personal.umich.edu/~cyoungk/schnittkebio.htm](http://personal.umich.edu/~cyoungk/schnittkebio.htm). Web. 5 June 2010.
- Koffka, K. *Principles of Gestalt Psychology*. Routledge, 1999. Print.
- Kraehenbuehl, D., and E. Coons. "Information as a Measure of the Experience of Music." *The Journal of Aesthetics and Art Criticism* 17.4 (1959): 510–522. Print.
- Kramer, J. D. "Moment Form in Twentieth Century Music." *The Musical Quarterly* 64.2 (1978): 177–194. Print.
- Lamont, Alexandra, and Nicola Dibben. "Motivic Structure and the Perception of Similarity." *Music Perception: An Interdisciplinary Journal* 18.3 (2001): 245–274. Print.
- Lansky, Paul, and George Perle. "Parameter." *Grove Music Online* n. pag. Print.
- Larkey, L. B, and A. B Markman. "Processes of Similarity Judgment." *Cognitive Science* 29.6 (2005): 1061–1076. Print.
- Lempel, A., and J. Ziv. "On the Complexity of Finite Sequences." *Information Theory, IEEE Transactions on* 22.1 (1976): 75–81. Print.
- Lewin, David. *Generalized Musical Intervals and Transformations*. Oxford University Press, 2007. Print.
- Longuet-Higgins, H. C., and C. S. Lee. "The Rhythmic Interpretation of Monophonic Music." *Music Perception* 1.4 (1984): 424–441. Print.
- Longuet-Higgins, H. Christopher, and Christopher S. Lee. "The Perception of Musical Rhythms." *Perception* 11.2 (1982): 115. Print.
- Losada, C. C. "Between Modernism and Postmodernism: Strands of Continuity in Collage Compositions by Rochberg, Berio, and Zimmermann." *Music Theory Spectrum* 31.1 (2009): 57–100. Print.
- Maus, F. E. "Concepts of Musical Unity." *Rethinking music* (1999): 171–192. Print.
- McAdams, S., and J. C Cunible. "Perception of Timbral Analogies." *Philosophical transactions: Biological sciences* (1992): 383–389. Print.
- McAdams, S., and D. Matzkin. "Similarity, Invariance, and Musical Variation." *Annals of*

- the New York Academy of Sciences* 930.1 (2001): 62–76. Print.
- Medin, Douglas L., Robert L. Goldstone, and Dedre Gentner. "Respects for Similarity." *Psychological Review* 100.2 (1993): 254–278.
- Messiaen, Olivier. "Oiseaux Exotiques." *Des canyons aux étoiles Oiseaux exotiques ; Couleurs de la cité céleste*. Cond. Esa-Pekka Salonen. Perf. Paul Crossley and London Sinfonietta. New York, N.Y.: CBS Records Masterworks, 1988. Audio Recording.
- Miller, G. A., and P. N. Johnson-Laird. *Language and Perception*. Cambridge university press Cambridge, 1976. Print.
- Nancarrow, Conlon. "Study No. 21." *Nancarrow: Studies for Player Piano [Box Set]*. Wergo Germany, 2000. Audio Recording.
- North, A. C, and D. J Hargreaves. "Subjective Complexity, Familiarity, and Liking for Popular Music." *Psychomusicology: Music, Mind and Brain* 14.1-2 (2010): 77–93. Print.
- Orpen, K. S, and D. Huron. "Measurement of Similarity in Music: A Quantitative Approach for Non-parametric Representations." *Computers in music research* 4 (1992): 1–44. Print.
- Palmer, C., and C. L. Krumhansl. "Mental Representations for Musical Meter." *Journal of Experimental Psychology: Human Perception and Performance* 16.4 (1990): 728–741. Print.
- Pan, D. "A Tutorial on MPEG / audio Compression." *Multimedia, IEEE* 2.2 (1995): 60–74. Print.
- "Parameter." *The New Oxford American Dictionary*. 2nd ed. 2005. Apple Application.
- Planets. "O People." *Planets*. Distile Records, 2007. Audio Recording.
- Pressing, J. "Cognitive Complexity and the Structure of Musical Patterns." *Proceedings of the 4th Conference of the Australasian Cognitive Science Society*. 1999. Print.
- "Productive Discontinuities." <http://www.havergalbrian.org/productive.htm>. Web. 5

- Mar. 2010.
- Randel, D. M. *The Harvard Dictionary of Music*. Belknap Press, 2003. Print.
- Rock, I. "The Description and Analysis of Object and Event Perception." *Handbook of Perception and Human Performance*. Vol. 2. 2 vols. New York: John Wiley and Sons, 1986. 33–1. Print.
- Rollins, Sonny et al. "My Reverie." *Tenor Madness*. Comp. Larry Clinton. Giants of Jazz ; Distributed by SAAR, 1990. Audio Recording.
- Savardi, Ugo, ed. *The Perception and Cognition of Contraries*. Milan, Italy: McGraw-Hill, 2009. Print.
- Shepard, R. N. "Stimulus and Response Generalization: A Stochastic Model Relating Generalization to Distance in Psychological Space." *Psychometrika* 22.4 (1957): 325–345. Print.
- . "The Analysis of Proximities: Multidimensional Scaling with an Unknown Distance Function. I." *Psychometrika* 27.2 (1962): 125–140. Print.
- . "The Analysis of Proximities: Multidimensional Scaling with an Unknown Distance Function. II." *Psychometrika* 27.3 (1962): 219–246. Print.
- Shmulevich, I. et al. "Perceptual Issues in Music Pattern Recognition: Complexity of Rhythm and Key Finding." *Computers and the Humanities* 35.1 (2001): 23–35. Print.
- Shmulevich, I., and D. Povel. "Measures of Temporal Pattern Complexity." *Journal of New Music Research* 29.1 (2000): 61. Print.
- Shmulevich, Ilya, and Dirk-Jan Povel. "Complexity Measures of Musical Rhythms." *Rhythm Perception and Production*. Lisse, The Netherlands: Swets & Zeitlinger Publishers, 2000. 239. Print.
- Smith, L. M, and H. Honing. "Evaluating and Extending Computational Models of Rhythmic Syncopation in Music." *Proceedings of the International Computer Music Conference*. Citeseer, 2006. 688–691. Print.
- Snyder, Bob. *Music and memory: an introduction*. Cambridge, Mass.: MIT Press, 2000.



- Print.
- Strauss, Johann. *Voices of Spring* (Waltz). Arr. Sebastian Spychaj. <http://www.free-scores.com/download-sheet-music.php?pdf=5935>. Web. 18 Jan. 2011.
- Tanguiane, Andranick S. *Artificial Perception and Music Recognition*. Springer-Verlag, 1993. Print.
- Toussaint, G. "A Mathematical Analysis of African, Brazilian, and Cuban Clave Rhythms." *Townson University*. CiteSeer, 2002. Print.
- Toussaint, G. T. "A Comparison of Rhythmic Similarity Measures." *Proc. 5th International Conference on Music Information Retrieval*. CiteSeer, 2004. 242–245. Print.
- Tversky, A. "Features of Similarity." *Psychological review* 84.4 (1977): 327. Print.
- Victor, J. D., and K. P. Purpura. "Metric-space Analysis of Spike Trains: Theory, Algorithms and Application." *Network: computation in neural systems* 8.2 (1997): 127–164. Print.
- Wessel, D. L. "Timbre Space as a Musical Control Structure." *Computer music journal* 3.2 (1979): 45–52. Print.
- Wiggins, G. A. "Models of Musical Similarity." *MUSICAE SCIENTIAE* 11.I (2007): 315. Print.
- Williams, Joseph M. *Style: Lessons in Clarity and Grace*. 9th ed. Pearson/Longman, 2006. Print.
- Zorn, John et al. *Naked city*. New York, N.Y.: Elektra/Nonesuch, 1989. Audio Recording.
- Zorn, John. "Road Runner." *Manhattan Cascade*. Perf. Klucevsek, Guy. Composers Recordings, 1994. Audio Recording.