

# Musical Qualia, Context, Time, and Emotion

Joseph A. Goguen

Department of Computer Science & Engineering  
University of California at San Diego, USA

**Abstract:** Nearly all listeners consider the subjective aspects of music, such as its emotional tone, to have primary importance. But contemporary philosophers often downplay, ignore, or even deny such aspects of experience. Moreover, traditional philosophies of music try to decontextualize it. Using music as an example, this paper explores the structure of qualitative experience, demonstrating that it is multi-layer emergent, non-compositional, enacted, and situation dependent, among other non-Cartesian properties. Our explanations draw on recent work in cognitive science, including blending, image schemas, and sensory memory, as well as on phenomenology. A hierarchical structure transformation based complexity theory is applied to obtain a non-linear dynamical systems explanation of qualia and emotion that respects phenomenological insights about time, including retention and protention. The complexity measure provides both a metric structure and a potential function, on spaces of pieces that are constructed using given elements and transformations, with weights that reflect their cognitive difficulty. However, the approach is not reductionist; using improvisation and the evolution of musical notation as data, we argue that situatedness, especially enactment and social context, are key aspects of musical consciousness.

## 1 Introduction

This paper addresses issues in the philosophy, cognitive science, and sociology of music, clustered around the notion of **qualia**<sup>1</sup>, which (roughly speaking) are the qualitative aspects of conscious experience. In particular, we consider the structure of qualia, that is, their parts, and how those parts combine. Other concerns include contextuality, memory, saliency, time and emotion.

Sections 2 and 3 focus on the philosophy of music, with Section 2.1 devoted to a very quick historical survey, and with Section 2.2 exploring some notions of context and situatedness, including the “fourth person” methodology that is used in this paper. Section 3 focuses on qualitative aspects of music, sketching some positions that philosophers have taken, including the separation of subjective from objective aspects, treating mainly the latter, and dubbing the former “qualia.” We propose an alternative approach, supported by examples<sup>2</sup> demonstrating some curious properties of musical qualia, such as having heterogeneous components with multiple levels, combining non-compositionally, and appropriating temporally prior qualia (thus violating some naive ideas about temporal linearity of perception). These phenomena are interesting in themselves, and they also contribute to skepticism about the ontological status of qualia. Section 3.1 discusses time consciousness, including Edmund Husserl’s notions of fresh memory, retention, and protention.

Section 4.1 discusses some cognitive science of music, arguing for the relevance of recent developments in cognitive linguistics, including image schemas and blending. Section 4.2 extends these ideas to handle structure more effectively, while Section 4.3, which is the technical heart of the paper, develops a hierarchical structural complexity theory with transformations, which is used to

---

<sup>1</sup>The singular form of this perhaps unfamiliar word is “quale.”

<sup>2</sup>The lecture version of this material featured short piano segments performed by Ryoko Goguen as illustrations.

model qualia, and then extended to construct a dynamical systems model of musical comprehension. Section 5 concerns the sociology of music, arguing that many phenomena discussed in previous sections, including qualia, have a significant social component. It takes as data an improvisation by Ryoko Goguen (which is documented in Appendix A), and some observations on the evolution of musical notation. Section 6 draws some conclusions about cognitivism, and about the origin, structure, and study of emotion and consciousness.

The study of music is usually the study of representations of music, mostly scores, in fact, mostly formal properties of eighteenth and nineteenth century scores in the European classical tradition. But scores are not music, and digital recording technology makes possible much more detailed representations of a much broader range of styles, including those for which scores do not exist and/or would not be very helpful even if they did exist. However, this still leaves us dealing with representations. Actually, I do not think it desirable to banish representations; the problem is with the “objective” third person approach usually taken to representations, rather than with their existence or their use. Our proposed alternative is the “fourth person” method sketched in Section 2.2, which requires *listening* to be taken much more seriously, including its social situatedness. We do this not only through close attention to our responses to music, but also in Section 4.3, by constructing *models* of listening, and comparing their responses with our own, with results from experimental psychology, and with relevant philosophical and musicological literature.

## Acknowledgements

I thank Dr Ken Mogi of Sony for his suggestion to write up my observations on musical qualia, Dr Volkmar Klien and Prof Dieter Kaufmann of the Performing Arts University of Vienna, and Prof Robert Trapp of the University of Vienna, for their hospitality during a visit in which some of these ideas were given in preliminary form, Prof Mitsu Okada for inviting a lecture on this subject at Keio University, Profs David Borgo and Prof George Lewis of UCSD, and Drs Erik Myin and Volkmar Klien for valuable comments and encouragement, Profs Chaya Czernowin and Chinary Ung of UCSD for putting up with a very early version in their class on metaphor and music, and Dr Jun Tani of the Brain Science Institute (Japan) for some stimulating conversations on dynamical systems and consciousness. Last but far from least, I thank my wife, Ryoko Amadee Goguen, for her help in performing this paper, and for immersing me in musical projects.

## 2 Some Philosophy of Music

Following a brief overview of some traditional Western philosophical views of music, some alternatives based on recent developments in sociology and cognitive science are discussed.

### 2.1 Some Traditional Western Views

The earliest influential views on music in the Western tradition come from Pythagoras, who claimed that music is essentially mathematical, based on his discovery that the ratios of the tones in musical intervals are simple rational numbers. Attempts to reduce music to mathematics are by no means confined to antiquity; for example, the great contemporary Greek composer Yannis Xenakis has written extensively (though not always coherently) in this mode. Returning to antiquity, Plato, in his *Republic* [64], wrote that “... music education is of paramount importance. More than anything else, rhythm and melody find their way straight to the deepest parts of our being,” which is consistent with what many musicians and educational psychologists think today. But Plato also seems aligned with many contemporary political conservatives, in wishing to “clean our republic of

depravity [in the arts] ... connected with vulgarity, insolence, madness, or other evils.” Thus, Plato was not a musical Platonist, believing that musical objects exist in some ideal realm, independently of humans; on the contrary, Plato’s view is focused on the moral character of musicians.

A common view during the Romantic era was that good music is an accurate representation of the emotions of its composer. Although still common among the general public, this view is rarely accepted by professionals<sup>3</sup>. Perhaps the earliest persuasive arguments against it are due to the Viennese critic Eduard Hanslick [45], who ascribed the emotional responses of listeners to analogies between dynamic patterns of music and emotions. Despite this apparently cognitivist view, Hanslick seems to have been a kind of Platonist, asserting that the beauty of music exists independently of any listener. Although Nelson Goodman is also interested in musical patterns, he identifies a musical work with its score and with the class of its authentic performances [43]. We argue against this “score nominalism” by noting that scores often do not exist, and for styles such as ambient music, would make little sense even if they did exist. More recent thinkers, e.g., Peter Kivy [51], take a cognitivist approach, defining music in terms of conscious patterns in the minds of listeners. Although music theory rather than philosophy, and written before the rise of cognitivism, the pioneering work of Leonard Meyer [56] takes a cognitive approach, defining musical experience instead of music, and emphasizing structure, anticipation, and emotion, in ways related to those in Section 4.3 below. Interestingly, Meyer’s three distinguishing characteristics of musical experience (on page 23 of [56]) do not mention sound, and (assuming the omission is deliberate) are so abstract as to encompass sequences of non-referential visual stimuli; I like this, and note its consistency with some postmodern views that music has an inherent multimedia character.

Theodor Adorno [1] takes a neo-Marxist social critical view of music, focusing on its roles in consumerism, class difference, etc. Despite this, one of his more philosophical statements, that “music is gesture,” embraces abstraction, and many recent composers, including Igor Stravinsky, support stronger views, in which music has only abstract musical content. Such views resemble those of modernist movements such as logical positivism, which attempts to reduce all “meaningful” information to abstract logic, as well as with psychological behaviorism, and with “eliminativist” trends which attempt to delete all mention of consciousness (see Section 3 below). Modernist Anglo-American analytic philosophy in general downplays or ignores the qualitative aspects of experience, despite strong experimental evidence of its importance, e.g., for motivation and recall [17]. Section 3.1 discusses some ideas from phenomenology, an approach that takes consciousness as central, and Section 6 includes a critique of cognitivism.

What seems clear from all this is that music is one of the least settled areas of philosophy. And of course Western views are not the only ones available; deep and interesting ideas can be found in many non-Western traditions, including those of China, Japan, and Bali. Moreover, even within the Western tradition, much of value can be learned from such often ignored genres as avant garde jazz, Delta blues, hip-hop, folk music, gospel, and heavy metal rock.

## 2.2 Context, Situatedness, and Embodiment

Traditional philosophical approaches like those discussed in the previous subsection tend to de-contextualize music. A common approach is to identify a musical object with some purely formal entity, which necessarily omits an enormous amount of relevant information. For example, no transcription, not even a spectral analysis, let alone a score, can capture all the nuances of an actual performance, which will include particular mediations by particular musicians, musical instruments, listeners, and rooms. The philosophical error that underlies such mistakes resembles

---

<sup>3</sup>Leonard Meyer [56] suggests that this is in part because emotional responses often become denatured through intellectual analysis and professional training.

that of mathematical Platonism, in which a transcendental ontological status is claimed for mathematical objects. Moreover, an emphasis on formal notation necessarily downgrades traditions, such as contemporary jazz, in which improvisation has great importance.

One approach to avoiding such problems is to reify the notion of context. But researchers in what we may call phenomenological sociology, for example, in ethnomethodology [69, 28], emphasize that context is dynamically emergent from activity, rather than fixed, definable in advance, formally representable, or separable from activity. Thus one should speak of *situated actions*, or *occasions of action*, rather than of *contextualized representations*, because neither situations nor their contexts are specifiable, representable, stable, or separable from their actual uses. For example, when a contemporary composer like Toru Takemitsu quotes Debussy, even if some measures are identical, the effect is radically different from the original; and similarly when Hans Werner Henze quotes Cuban folk songs or Monteverdi. Even in a Mozart piano sonata, the second entrance of a theme can have a very different effect from the first.

Paul Dourish [23] gives an insightful discussion of context in connection with current trends towards ubiquitous or “context aware” computing; the problem addressed by this field is how to use powerful new sensor technologies to make computational systems more responsive to their users’ physical and social settings, as those users move through and modify these settings. This has turned out to be unexpectedly difficult, and Dourish claims this is essentially for reasons like those described above.

Musical Platonism, score nominalism, cognitivism, and modernist approaches in general, all assume the primacy of representation, and hence all flounder for similar reasons. Context is crucial to interpretation, but it is determined as part of the process of interpretation, not independently or in advance of it. Certain elements are recognized as the context of what is being interpreted, while others become part of the emergent musical “object” itself, and still others are deemed irrelevant. Moreover, the elements involved and their status can change very rapidly. Thus, every performance is uniquely situated, for both performers and listeners, in what may be very different ways. In particular, every performance is **embodied**, in the sense that very particular aspects of each participant are deeply implicated in the processes of interpretation, potentially including their auditory capabilities, clothing, companions, musical skills, prior musical experiences, implicit social beliefs (e.g., that opera is high status, or that punk challenges mainstream values), spatial location, etc., and certainly not excluding their reasons for being there at all (this is consistent with the cultural historical approach of Lev Vygotsky [77]). More than that, participants **enact**, or actively construct, the context of musical experience, sometimes quite visibly so, as in rock concerts and karaoke parties, e.g., with shouts of encouragement or disparagement. David Sudnow gives an insightful ethnomethodological description of learning to play jazz piano, emphasizing the role of the body (hence its title, *Ways of the Hand* [70]).

Most scientific studies of art are problematic for similar reasons. In particular, the third person, objective perspective of science requires a stable collection of “objects” to be used as “data,” which therefore become decontextualized, with their situatedness, embodiment, and interactive social nature ignored. Moreover, any choice of a fixed dataset necessarily makes some arbitrary presuppositions about the nature of what is analyzed. For example, typical textbooks on harmony deal exclusively with traditional Western classical music, e.g., [2], as do typical artificial intelligence projects, e.g., Gerhard Widmer [79] uses electronic transcriptions of performances of Mozart piano sonatas as data. In the visual arts, Ramachandran and Hirstein [65] use optical illusions and classical Indian erotic art as data, while Zeki uses mainly Renaissance and early modern art [82], and Taylor classifies the drip art of Jackson Pollock using fractal dimension [72]. Although certainly valuable in their own domains, one should be cautious about generalizing the conclusions of such studies, particularly for contemporary art, a major theme of which has been to question what art

is, and to explore its boundaries [35].

This paper draws on data from both science and phenomenology, in a spirit similar the “neurophenomenology” of Francisco Varela [74], as a way to reconcile first and third person perspectives, by allowing each to impose constraints upon the other. Such approaches acknowledge that the first and third person perspectives reveal two very different domains, neither of which can be reduced to the other, but they also deny that these domains are incompatible. It is important to note that phenomenology is *not* merely a more disciplined form of introspection, but rather attempts to enhance sensitivity to ordinary experience, without distortion caused by separation from its natural context<sup>4</sup>. The approach of this paper goes beyond that of Varela in the following ways: (1) it admits not just neuroscience, but also the third person perspectives of other sciences, especially evolution, as well as of critical studies in the arts and humanities; (2) it takes account of the “second person” perspective of society, in the style illustrated by Section 5; and (3) it places greater emphasis on the lived experience and personal transformation of scientists, technologists, critics, etc.<sup>5</sup>. The intention of this approach is to advance scientific knowledge, technical practice, personal awareness, and the health of society. I propose to call it the **fourth person method** or **perspective**, since it encompasses the first, second, and third person perspectives, and because of an analogy with the “fourth moment,” which is described at the end of Section 6. This paper illustrates the application of this approach to music, in which all three perspectives play an important role. It seems to me that such an approach is quite natural, since it is what many of us do anyway, and hence it does not require any elaborate methodological justification. However, the fact that it challenges some deeply entrenched strictures of conventional science, does raise issues that require further discussion, although this paper is not the right place to do it.

### 3 Musical Qualia

Listeners generally consider the qualitative feel of music to be its essence and its chief attraction, emotion of course being a prime example. But as noted in the previous section, many philosophers ignore the qualitative aspect of experience, and some even deny its existence. When it is admitted that experience is not just perceptual and conceptual, a common approach is to “bracket” or exile the qualitative aspects, and concentrate attention on aspects that are reducible to scientific analysis, which for music would include duration in milliseconds, intensity in decibels, and spectral analysis of timbres; the recalcitrant residua are then dubbed “qualia” and thereafter largely ignored. Under this view, qualia are what’s left after the objective aspects are subtracted (a variant is that qualia are what’s left after intentionality is subtracted – see [14] for a relatively clear exposition of this and related issues). The resulting residual category<sup>6</sup> could include non-conceptual associations, emotional content, etc. Examples from perception are often given, since what we perceive often differs significantly from what physical instruments reveal as objective. Perhaps the most commonly used example is the “redness of red”; some others are the special qualities of a slowly bent major to minor third, the smell of onion, and the blueness of sky. Of course any piece of music is an example, although the philosophical literature tends to shy away from examples having strong emotional connotations.

---

<sup>4</sup>Since training in phenomenological observation is difficult to obtain, Buddhist meditation is often mentioned as an alternative, which moreover has a good track record in anticipating results from cognitive science (see Section 6).

<sup>5</sup>Our method also differs from the “heterophenomenology” of Dennett [20], though this discussion lies outside the scope of this paper.

<sup>6</sup>Residual categories often have a special sociological flavor, as “the other,” which is partially unknown, partially feared, partially valorized, and always ambiguous. Examples less esoteric than qualia include racial and ethnic minorities, new diseases, and in some circles, consciousness itself.

A difficulty with qualia advocacy is that it tends to reify qualia, giving them independent existence as Platonic entities, and introducing a fundamental ontological distinction between subjective and objective aspects of experience. For example, David Chalmers [11] claims there are two fundamental “world substances,” matter and information. It seems to me that such approaches are what Daniel Dennett [19] argues against, although he is sometimes accused of arguing against the existence of qualitative experience as such. A genuine eliminativist is Francis Crick (of DNA fame), who wrote “You’re nothing but a pack of neurons” [15]; see also [12]. One motivation for qualia is to challenge dominant reductionist paradigms of neuroscience, analytic philosophy, and experimental psychology. For example, Chalmers [11] poses what he calls “the hard problem,” which is to explain qualia in the language of the hard sciences, his presumption being that this may not be possible. We argue that cognitive and qualitative aspects of experience are inseparable, even though first and third person approaches artificially separate them.

Whatever approach is taken, qualia are often considered to be *atomic*, i.e., non-reducible, or without constituent parts, in harmony with doctrines of logical positivism, e.g., as often attributed to Wittgenstein’s *Tractatus* [80]. Though I have never seen it stated quite so baldly, the theory (but perhaps “belief” is a better term, since it is so often implicit) seems to be that qualia atoms are completely independent from elements of perception and cognition, but somehow combine with them to give molecules of experience. For example, this view can be inferred from standard musical notation, where the qualitative, or “expressive,” aspects are given in natural language, typically Italian (e.g., “*dolorisimo*,” “*fantastico*,” “*con spirito*”), whereas pitch and duration are given in an abstract symbolic language<sup>7</sup>. Interestingly, tempo occupies an intermediate position, with both kinds of notation in active use, though musicians often interpret even quite precise tempo markings (in beats per minute) as if they were qualitative. Moreover, volume (intensity) markings are also intermediate, the notation used consisting of contractions of natural language (again from Italian) rather than precise decibel levels. Contemporary classical music elaborates these conventions in numerous ways, with varying degrees of success, usually to give composers more control over performers; for example, the Italian contemporary classical composer Luigi Nono distinguishes numerous kinds of fermata<sup>8</sup> in some of his scores. It is easy to find examples of non-atomic musical qualia, having a feeling tone that are not simple combinations of the feeling tones of their constituents; this means that qualia are non-compositional. For example, a minor triad has a feeling that cannot be inferred by listening to its three notes on three sufficiently separated occasions.

The view of this paper, which is influenced by phenomenology, especially that of Edmund Husserl [49] and Martin Heidegger [47], is that qualia are not separable from experience, and that experience is primary. Given this view, it is convenient to use the word “quale” to emphasize that we address this primary experiential dimension, rather than objective measurements, and to informally define “**qualia**” as conscious experiences having unity<sup>9</sup> and duration. This definition does not separate “subjective” aspects from some other, allegedly “objective” aspects: we could perhaps say that it views everything as subjective, but it would be much better to say that we consider the distinction meaningless at this level of analysis. This approach to qualia escapes criticisms like those advanced by Dennett [19]. Actually, our view of the nature of qualia is such that it is more appropriate to *model* the production of qualia than it is to define qualia; see the developments in Section 4.3.

Musical practice confirms that separation of quantitative and qualitative aspects is illusory, a matter of notational convenience and tradition, not an inherent property of music itself. For

---

<sup>7</sup>But this notation is somewhat iconic, in the technical sense of Charles Sanders Peirce [63].

<sup>8</sup>These are rests of free duration.

<sup>9</sup>In the sense of being experienced as whole, even though possibly also experienced as having distinguishable parts.

example, the same written phrase is performed in a very different way if it occurs in a Mozart score or a Count Basie score, and experienced listeners can easily detect many variants even within a single genre, as well as very many genres, realized for example through small variations in timing (e.g., Kansas City swing, Mersey beat, and Afro-Cuban bebop).

Here are some further curious properties of musical qualia: Small changes can have huge effects, e.g., changing one note of a traditional diatonic melody by a quarter tone. Also, the same quale may have components of many different kinds; for example, the uplifted feeling of the final chorus in the Gloria of Bach's *B Minor Mass* arises from a very complex interaction of melody, harmony, voicing/orchestration, tempo changes, key changes, and dynamics<sup>10</sup>. Moreover, temporally prior qualia may be appropriated (depending on saliency), thus violating overly simple ideas about the temporal linearity of perception; for example, in a rapid chromatic run, most notes will lose their identity, except (potentially) the first and last. Ambiguities very possible, and are often exploited by composers, e.g., in false endings. There are also important large grain effects, such as similarity and contrast in the development section of a classical sonata. A particularly interesting phenomenon is when some material "jumps" from being a sequence to being the operation of a process; my favorite example occurs in the third piece of Anton Webern's *Five Pieces for Orchestra*, opus 10, when a sequence of irregular percussive sounds is suddenly perceived as the erratic operation of some malfunctioning device, such as a broken clock; however, qualia within the cycles still retain individuality. Such phenomena demonstrate that qualia are far from being atomic or compositional, and hence far from being Cartesian, in the sense of Descartes' *Discourse on Method* [21], which recommends analyzing phenomena into "as many parts as possible," each of which is "clear and distinct," and then "gradually and by degrees reaching towards more complex knowledge"<sup>11</sup>.

Musical qualia provide concrete counter-examples to certain philosophical positions, including Brentano's thesis on intentionality, which (under some interpretations, e.g., [14]) holds that every thought has an intentional component, some "thing" that it is about. For, although some musical thoughts are "about" something (e.g., Roman fountains), most are not, unless perhaps a very broad notion of "about" is allowed, which includes intra-musical relationships, since most musical thoughts are only about other aspects of music. In addition, the end of Section 3.1 refutes some of Hume's views on time [48]. Section 6 discusses some further disconfirmations and implications.

### 3.1 Time Consciousness

The subjective perception of time, which is clearly crucial for music, has been investigated deeply by phenomenologists, especially Edmund Husserl, whose work makes explicit use of music, at least as a metaphor<sup>12</sup> [49]. Husserl describes some important ways in which our actual subjective experience of the three-fold division into past, present, and future differs from the usual physical scientific analysis. He observes that the flow of time is a continual "sinking away" into the past, and that experience can only occur in the present. Thus, the past can be experienced in the present, but in a different *mode* from that in which it was originally experienced. Husserl names this mode **retention**; it allows us to distinguish present experience from past experience, even though both are experienced in the now. In addition, while listening to music, we are constantly anticipating what may come next, a mode that Husserl names **protention**; this is what makes it possible for us

---

<sup>10</sup>In music, "dynamics" usually refers mainly to amplitude, even though all parameters are constantly changing.

<sup>11</sup>Although Descartes did not treat qualia as such, nor did he name his smallest parts atoms, it seems fair to say that his clear description of his method, and his very impressive application of it in reducing geometry to arithmetic with Cartesian coordinates, were enormously influential for all subsequent philosophy of science, which often impose a stricter interpretation than was probably intended by Descartes.

<sup>12</sup>Given the methods of phenomenology, it is likely that Husserl actually used music in these investigations.

to experience qualia of surprise. Note that protention is not symmetrical with retention: it does not provide a single “image” of what is coming, but rather a complex network of expectations about what might come.

Husserl’s observations on retention are confirmed by recent research on temporal cognition of music [66], identifying a pre-conscious buffer of about 10 seconds. This buffer corresponds to what Husserl calls **fresh memory**<sup>13</sup>, and it differs from the more familiar short term memory and long term memory, in that it is not conscious, i.e., it works whether or not we are aware of it (though conscious attention may make it work better). Of course, the traditional short and long term memories also play important roles in music, allowing us to recognize themes that recur in a piece, and pieces that we have previously heard. Note that reification through repetition of a sequence to a process works best if the sequence is short enough to fit into the fresh memory buffer, especially if there are continual variations, as in the third piece of Webern’s opus 10. There do not yet appear to be any neuro-cognitive studies of protention.

For Husserl, “objective moments” of time are not pre-determined, but rather, objects-in-time arise through processes of retention and protention. His goal was to discover the origin of time, and these processes are his answer. Our observations on musical qualia demonstrate that only certain salient configurations of objects and relations become so solidified, and that it is this solidification that creates apparent “moments” in time; other less salient configurations are less solidified, and less likely to be retained or “objectified” as events. Moreover, these moments, unlike those of physics, have duration, or are “temporally thick,” since they relate to real events which take time to process (the processing time is about one tenth of a second, but varies considerably with conditions). In addition, relationships that hold a salient configuration together are retained in memory with their salient constituents, and their qualia. Indeed, there is much experimental evidence that qualia, at least emotional qualia, function as “indices” for the retrieval of memories, as well as playing important roles in many other mental processes, including reasoning, e.g., see [16, 17]. Husserl does not discuss qualia, but our musical observations shows that only salient perceptions become phenomenological events, and that only these have an associated qualia; we will argue in Section 4.3 that this has to do with protention and anticipation.

These observations also disconfirm aspects of David Hume’s pointilist theory of time [48]. Hume was right that time consists of discrete episodes, but wrong that these are discrete point-like instants, since musical qualia are clearly “temporally thick,” i.e., have durations that cannot meaningfully be broken into points; moreover, it is easy to find musical examples in which an event has ill-defined or gradual onset and/or conclusion (e.g., in group improvisations by the Art Ensemble of Chicago). Alfred North Whitehead [78] is another philosopher in this category. His notions of prehension and concrescence, and his “epochal” theory of time, have similarities with Husserl and with the theory of qualia in this paper. Nor should we fail to mention William James [50], whose present ideas are further briefly discussed in Section 6.

Jun Tani [71] observed that during a sensory-motor learning task, a simulated robot with a certain hierarchical neural net architecture exhibited transitions between phases of coherence and incoherence in its anticipation of perceptions, where the resulting coherent chunks correspond to the structure of the task. Tani relates this to Husserl’s phenomenology of time, in that the chunks are “immanent” for the robot. Varela [75] also studied some connections between cognitive neuroscience and Husserl’s phenomenology of time, in part as an illustration of the neurophenomenological method introduced in [74].

---

<sup>13</sup>A usual term for it in psychology is “sensory memory.” Note also that short term auditory memory is about 30 seconds (unless refreshed by repetition), and thus is quite distinct from sensory memory.

## 4 New Models for Understanding Music

This section suggests that we can better understand what it means to understand music by constructing models of *how* we understand music. Like all models, ours will be partial descriptions constructed for particular purposes, but also, since they are mathematical models, they can be more precise, and more falsifiable, than more traditional philosophical or musicological theories. Our approach is intended to apply to contemporary musical manifestations, such as noise music, digital multimedia productions, free jazz improvisation, music sculptures, environmental music, etc. Though more rigorous, the theory has many points of contact with the innovative work of Leonard Meyer [56], particularly regarding the importance of anticipation in music, and can even be seen as an attempt to update Meyer’s work; however, our approach also differs in some significant ways, particularly its technical machinery for handling structure and hierarchical complexity. There are of course many other mathematical approaches to music, e.g., see the collection [5], though usually they are limited to classical music, and often have a Platonist philosophy.

### 4.1 Metaphor and Blending

Research in cognitive linguistics by George Lakoff and others under the banner of “conceptual metaphor theory” (abbreviated “CMT”) has greatly deepened our understanding of metaphor [54, 53], showing that many metaphors come in families, called **image schemas**, that share a common pattern. One example is BETTER IS UP, as in “I’m feeling up today,” or “He’s moving up into management,” or “His goals are higher than that.” Some image schemas, including this one, are grounded in the human body<sup>14</sup> and are called **basic image schemas**; they tend to yield the most persuasive metaphors. Such image schemas do occur in music; for example, an angelic choir with high voices instantiates the BETTER IS UP schema.

Fauconnier and Turner [26, 27] have studied **blending**, or **conceptual integration**, claiming it is a basic human cognitive operation, invisible and effortless, but nonetheless fundamental and pervasive, appearing in the construction and understanding of metaphors, as well as in many other cognitive phenomena, including grammar and reasoning. Many simple examples are blends of two words, such “houseboat,” “roadkill,” “jazz piano,” “computer virus,” “classical composer,” and “melodramatic conductor.” To explain such phenomena, blending theory (abbreviated “BT”) posits that concepts come in clusters, called **conceptual spaces**, which consist of certain items and certain relations that hold among them. Such spaces are relatively small, transitory constructs, selected on the fly from larger domains, to meet an immediate need, such as understanding a particular phrase or sentence. However, we do not assume that they are necessarily the **minimal** such spaces needed to understand a given blend, since that can only be determined after the blend has been understood. Moreover, different blends may omit different elements of the input spaces, and it may also be necessary to recruit additional information from other spaces in order to understand a blend. The abstract mathematical structure of a conceptual space consists of a set of atomic elements together with a set of relation instances among those elements [34]; of course, such a representation necessarily omits the qualitative, experiential aspects of what is represented (the qualia). **Conceptual mappings** are partial functions from the item and relation instances of one space to those of another, and **conceptual integration networks** are networks of conceptual spaces and mappings that are to be blended together.

---

<sup>14</sup>The source UP is grounded in our experience of gravity, and the schema itself is grounded in everyday experiences, such as that when there is more beer in a glass, or more peanuts in a pile, the level goes up, and that this is a state we often prefer; therefore the image schema MORE IS UP, discussed in [53], is even more basic.

The simplest blends<sup>15</sup> have the form of Figure 1, where  $I_1$  and  $I_2$  are called the **input spaces**,  $B$  is called the **blend space**, and  $G$  the **generic space**; the latter contains conceptual structure that is shared by the two input spaces<sup>16</sup>. A **blendoid** of  $I_1, I_2$  **over**  $G$  consists of a space  $B$  together with conceptual mappings  $I_1 \rightarrow B$ ,  $I_2 \rightarrow B$ , and  $G \rightarrow B$ . There may be many such blendoids, but relatively few are likely to be interesting. Therefore additional principles are needed for identifying the most interesting possibilities, so that we can define a **blend** to be a blendoid that is *optimal* with respect to these principles. Fauconnier and Turner suggest a number of “optimality principles” that serve this purpose (see Chapter 16 of [27]), but they are too vague to be easily formalized. A tentative and difficult but precise mathematical approach is given in Appendix B of [40], based on a modification of the category theoretic notion of “pushout” [55]; this modification takes advantage of an ordering relation on morphisms, with respect to their quality, as discussed in [34]. The intuition is that nothing can be added to or subtracted from such an optimal blendoid without violating consistency or simplicity in some way. However, there can still be more than one blend in this sense, as an example discussed below will make very clear. It should also be noted that this notion of blend easily generalizes to any number of semiotic spaces, and even to arbitrary diagrams of semiotic spaces and morphisms, for which there are many significant applications. Thus, the emphasis on double scope blending in [27] seems somewhat out of place in algebraic semiotics, because its major applications typically involve multiple “scopes” arising from multiple spaces and morphisms among them.

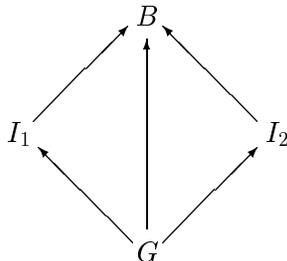


Figure 1: A Blend Diagram

---

It has perhaps not been sufficiently emphasized in the BT literature that blending does not always give a unique result. For example, the following are four different blends of conceptual spaces for “house” and “boat”: houseboat; boathouse; amphibious RV; and boat for moving houses. The last may be a bit surprising, but I once saw such a boat in Oban, Scotland, transporting prefabricated homes to a nearby island. There are also some other, even less obvious blends [38].

The CMT view of metaphor associates aspects of one domain to another, and describes this association using a mapping, of which the target domain concerns what the metaphor is “about.” On the other hand, BT views metaphors as “cross-space mappings” that arise from blending conceptual spaces extracted from the domains involved. For example, the metaphor “my love is a rose” arises from blending conceptual spaces for “my love” and “rose,” such that the identification of

---

<sup>15</sup>This diagram is “upside down” from that used by Fauconnier and Turner, in that our arrows go up, with the generic  $G$  on the bottom, and the blend  $B$  on the top. This is due to a pervasive and natural duality between theories and models, in the sense that these terms are used in mathematical logic; see Section 4.2 for further discussion. Our convention is also consistent with the way that such diagrams are usually drawn in mathematics, as well as with the image schema MORE IS UP (since  $B$  is “more”). Also, Fauconnier and Turner do not include the map  $G \rightarrow B$ .

<sup>16</sup>However, [34] uses the term “**base space**”, because it is more descriptive of how this space is used in applications to user interface design.

the two items “love” and “rose” in the blend space gives rise to a correspondence between certain items in the rose space and the target love space. Such **metaphoric blends** are *asymmetric*, in that as much as possible of the target space is imported into the blend space, whereas only key aspects from the source space, associated with elements that have been identified with elements of the target space, are imported, e.g., sweet smell and attractive color; moreover, names from the top space take precedence over those in the source space, so that relations in the source space become “attributed” to items in the target space. Our approach differs from orthodox BT not only in that we allow many more kinds of structure in our spaces (as discussed below in Section 4.2), but also in that we do not first construct a minimal image in the blend space and then “project” that same material back to the target space, but instead, we build the entire result structure in the blend space. Thus it is not the case for us that, in forming the blend, elements are preferentially omitted from the target space, only to be restored upon projection, as with the procedure described in [44]. Since CMT has been mainly concerned with families of metaphors having a shared pattern, whereas BT has been more concerned with how novel metaphors can be understood, the two theories are compatible, and can both play a role in understanding complex language. This and related issues are discussed with many interesting details in [44].

The conceptual spaces, mappings and blending of cognitive linguistics seem well adapted for treating many aspects of literature, as in [73], as well as some recent trends in art, including (the very aptly named) conceptual art movement, and with the conceptual aspects of works in many other styles, which are often designed to provoke conceptual conflicts or to force unusual conceptual blends. One important application is the combination of music and lyrics, as skillfully studied using cross-domain mappings by Lawrence Zbikowski [81]. However, the framework seems too restricted for studying blending *within* music, e.g., harmony, polyphony, polyrhythm, etc., because musical structure is inherently hierarchical, and hence cannot be adequately described using only atomic elements and relation instances among them. Understanding how a particular melody, chord sequence, and rhythm can work together requires close attention to the component notes, phrases, chords and beats, as well as to their subcomponents. Fortunately, it appears that the added generality of semiotic spaces and semiotic morphisms, as sketched in the next subsection, is adequate for such purposes. In the UCSD Meaning and Computation Lab, Fox Harrell and I have an experimental blending algorithm, which has generated novel metaphors used, which in turn were used in generating poems [38].

## 4.2 Semiotic Spaces and Structural Blending

Before introducing algebraic semiotics and structural blending, it is good to be clear about their philosophical orientation. The reason for taking special care with this is that, in Western culture, mathematical formalisms are often given a status beyond what they deserve. For example, Euclid wrote, “The laws of nature are but the mathematical thoughts of God.” Similarly, the “situations” in the situation semantics of Barwise and Perry, which resemble conceptual spaces (but are more sophisticated – perhaps *too* sophisticated), are considered to be actually existing, ideal Platonic entities [7]. Somewhat less grandly, one might consider that conceptual spaces somehow already exist in the brain. However, the point of view of this paper is that all formalisms are constructed in the course of some task, such as scientific study or engineering design, for the heuristic purpose of facilitating consideration of certain issues in that task. Under this view, all theories are situated social entities, mathematical theories no less than others; of course, this does not mean that they are not useful.

Classical semiotics was founded by Charles Sanders Peirce [63] and Ferdinand de Saussure [67] in the late nineteenth century. Peirce was an American logician concerned with problems of

meaning and reference, who concluded that these are relational rather than denotational, and who also made an influential distinction among modes of reference, as symbolic, indexical, or iconic (see footnote 7). Saussure, a Swiss linguist, wanted to understand how features of languages relate to meanings, and he emphasized binary features and denotational meaning. More recent thinkers, such as French literary theorist Roland Barthes [6], combined and extended these theories, creating a powerful language for cultural and media studies, which in various versions has been called semiotics, semiology, structuralism, and finally post-structuralism. However, this body of theory lacks the mathematical precision needed for scientific analysis, and also does not address dynamic signs, social issues such as arise in collaboration, or the systematic mapping of signs in one system to signs in another; moreover, many versions tend towards a Platonist view of signs. Algebraic semiotics, which was originally developed as a foundation for user interface design, attempts to overcome these problems. We now give an intuitive introduction to the way that it treats the structure of complex signs, such as musical scores, books, and graphical user interfaces. Details omitted here can be found in [34, 37, 36]; this theory originated in an early experimental study of multimedia learning [39], and was adapted to user interface design in [32].

A **semiotic system** or **semiotic theory** consists of: a **signature**, which gives names for sorts<sup>17</sup>, subsorts, and operations; some **axioms**; a **level ordering** on sorts having a maximum element called the **top** sort; and a **priority ordering** on the constructors at each level, where constructors are operations that build new signs from given parts. Sorts classify the parts of signs, among which data sorts provide values for attributes of signs (such as color and size). Axioms are constraints on the possible signs of a system. Levels express the whole-part hierarchy of complex signs, whereas priorities express the relative importance of constructors and their arguments; social issues play an important role in determining these orderings. This approach has a rich mathematical foundation, e.g., [40, 42, 41], since a signature plus equational axioms is an algebraic theory, on which there is a large literature. A theory-based approach is preferable to a more concrete set-based approach, because it allows both multiple models and open structure, both of which are important for applications. The first means, for example, there are many ways to play a score, while the second means that additional structure (such as another movement) or constraints (e.g., on tempo) can be imposed at a later time. In addition, it is more natural to treat levels and priorities in theories. Conceptual spaces correspond to the very special case of semiotic theories where there is only one sort, there are no operations except those representing atomic elements and relations, and axioms only assert that a relation holds of certain constants.

Associated with any theory are all the structures (called **models** or **algebras**) that provide *interpretations* for the things in the signature of the theory: sorts are interpreted as sets; operation symbols are interpreted as functions on these sets, with constant symbols interpreted as elements; all in such a way as to satisfy the axioms in the theory. The theory is thus a formal language for talking about such models. For example, the space of models for a theory of books consists of all books having the structure specified in the theory. Models and theories are *dual* ways of looking at the same thing, since associated with any class of structures is a unique most restrictive theory that has the given ones among its models.

Mappings between structures became increasingly important in twentieth century mathematics and its applications; examples include linear transformations (and their representations as matrices), continuous maps of spaces, differentiable and analytic functions, group homomorphisms, and much more. Mappings between sign systems are only now appearing in semiotics, as uniform representations for signs in a source space by signs in a target space. Since we formalize sign systems as algebraic theories with additional structure, we should formalize **semiotic morphisms**

---

<sup>17</sup>The word “sort” is used to avoid the ambiguities of the word “type.”

as mappings of theories that preserve the additional structure; however, these mappings must be partial, because in general, not all of the sorts, constructors, etc. are preserved in real examples. For example, the semiotic morphism from the rose space to the blend space for the metaphor “My love is a rose” (most likely) omits fertilizer and insects, while (possibly) preserving at least one of perfume and thorns. In addition to the structure of algebraic theories, semiotic morphisms should also (partially) preserve the priorities and levels of the source space. The extent to which a morphism preserves the various features of semiotic theories is an important determinant of its quality [40]. Semiotic morphisms can also be used to relate music to extra-musical elements. For example, letters of an alphabet can be associated with notes, so that sequences of notes can spell names, about which some information could be given, e.g., that “BACH” is the name of a composer who likes such tricks.

The simple form of blend in Figure 1 applies just as well to semiotic spaces and semiotic morphisms. In this generalization, the diagram is called a **structural blend diagram** and  $B$  a **structural blend**. Moreover, this also extends to diagrams with any number of semiotic spaces, and even to arbitrary diagrams of semiotic spaces and morphisms, in which case one may speak of a **structural integration network**.

### 4.3 Structural Complexity, Protention and Qualia

This section develops a model for the understanding of music, based on a complexity theory for hierarchical structures that are constructed through sharing and transformation<sup>18</sup>. The resulting theory appears to capture many features of musical qualia, as well as of Husserl’s notions of retention and protention. As before, readers should be aware that applications of such a formalization should be grounded in social, cognitive, and embodied reality, that no special ontological status is claimed for the abstractions involved, and that it is not intended as a grand general solution to all possible problems in the philosophy, psychology, sociology, etc. of music, let alone of consciousness.

Traditional Shannon information theory [68] has played a modest role in music theory, but has been rightly criticized for its inability to move beyond local features (such as so called  $n$ -grams) to larger grain structures, such as sonata form. More fundamentally, a theory that is based on probability necessarily makes some very dubious assumptions about the nature of music, such as that there are discrete atomic events, having fixed probabilities<sup>19</sup>. Complexity based information theories, along lines pioneered by Andrei Kolmogorov [52], are more attractive, since they can potentially encompass hierarchical structures, i.e., nested whole/part hierarchies. However, they seem unable to take account of features of memory like those discussed in Section 3.1.

A key concept of [31] is that of an analysis or “understanding” of a system  $S$  drawn from some given family or “space” of systems. Such an **analysis**  $A$  is a network of components, which yield  $S$  when combined, where each component in  $A$  is either atomic, a combination of other lower level components in  $A$ , or else a transformation of some lower level component in  $A$ , where the atomic components, the modes of combination, and the allowed transformations are fixed for the given family of systems, each with a numerical “weight” which reflects its cognitive difficulty. For music, the allowed components might be represented as functions from a temporal domain into a domain of musical notation; the set of such functions forms a “space” called the **state**

---

<sup>18</sup>The theory was introduced in [31], but research has advanced beyond what is described there, and moreover, some presuppositions of that paper no longer seem acceptable.

<sup>19</sup>Though this is not, of course, to say that probability cannot be used as a compositional technique, e.g., John Cage, Yannis Xenakis, Gyorgi Ligeti, and others have done so; but it is doubtful that most listeners attempt to infer those techniques from what they hear, either consciously or unconsciously.

**space**<sup>20</sup>. The **complexity**  $C(A)$  of an analysis  $A$  is the weighted sum of the complexities of the components of  $A$ , which in turn are also such weighted sums, and so on hierarchically, down to some atomic<sup>21</sup> elements, noting that reuse of an existing component or subcomponent will in general be weighted much less heavily than the first use. These weights are not probabilities, not even so called “subjective” probabilities, because cognitive difficulty is only partially determined by prior experience. Finally, the **structural complexity** of  $S$  is the non-negative real number denoted  $H(S)$ , which is the minimum of the complexities  $C(A)$  of all analyses  $A$  of  $S$ . The result of analysis is not just this number, but more significantly, yields a simplest hierarchical structural analysis  $A$  of  $S$  (though this analysis need not be unique); this analysis reveals not only small grain structures, but also large grain structures, as well as how all these structures are inter-related. Note that  $H(S)$  is *not* a measure of aesthetic preference, like that of George David Birkhoff [8].

Two unusual features of this theory are that components and transformations have weights, and that some components may be transformations of other components. In its application to music, the weight of a transformation should be determined by its cognitive difficulty. Thus, for classical music, a repeat of a (not too unwieldy) segment should have a very low additional complexity, and a transposition of it by a fifth or a fourth should have very little more, whereas more cognitively difficult transformations, such as retrograde inversion, should have a greater weight. Note also that this theory is not limited to standard musical notation; on the contrary, entirely different dimensions could be notated for a particular performance, such as musician biometrics (e.g., heart rate and posture), or audience acoustic energy, and of course spectral analyses or other less detailed analyses of timbre.

We are working to implement an algorithm to compute the minimum complexity and associated simplest structural analyses of temporal sequences, based on techniques of dynamic and backtrack programming. We found an early prototype to be very inefficient in both time and space. However, since human music listeners, as an aid to understanding music in real time, try to anticipate what might come next, based on what they have already heard, and on culture-specific conventions, it is natural to incorporate such features into a modified model, and implement it for experimentation. Let us call this model the **Anticipatory Model**. It will have a memory hierarchy, with sensory, short term, and long term components, where the latter includes templates for culturally preferred structures. The associated weights should correspond to the ease of retrieving information from the various layers of memory. For items in the sensory memory buffer, these weights should be very low, while items in short and long term memory should have greater weight. Although these modifications would lead to an algorithm that is much more efficient and closer to human listening, the resulting structural analyses would be unlikely to have minimum complexity; however, they should be close to minimum when the parameters and templates are a good match to a style of music. Methods for specifying the state space of possible hierarchical descriptions of music are described in [31], and are a special case of the semiotic systems described in Section 4.2. A subtle technical point that we only mention here in passing, is to reconcile the set-based sheaf models of [31] with the theory-based algebraic semiotics of Section 4.2; hidden algebra [36] can be used for this purpose.

The Anticipatory Model can be related to the phenomenology of music listening through several hypotheses: The first issue, which we call the **Main Hypothesis**, is that our hierarchical analyses, using psychologically and culturally appropriate components and weights, are understandings of works for appropriate audiences, and that a minimum complexity analysis gives a “best” such

---

<sup>20</sup>The space can be constrained by axioms that disallow some functions; also, multiple performers (or parts) can be handled by letting  $S$  take values in a product space of notational choices.

<sup>21</sup>These need not be atomic in the sense of having no component parts, but only in the sense of being treated as a whole; for example, the opening four notes of Beethoven’s fifth symphony are atomic in this sense for many listeners.

understanding, along with a precise structural analysis of the work in question. Second, the **Qualia Hypothesis** says that the weight of a unit corresponds (inversely) to its saliency, and the saliency of a unit gives its strength as a quale. This can explain some of the peculiar phenomenology of musical qualia discussed in previous sections. For example, it explains how the qualia of subphrases can be absorbed into the qualia of larger phrases: this occurs when new units of low weight are formed by incorporating prior material of originally higher weight. **Hypothesis S** says that a large-grain unit of low weight is one that is expected and remembered; a unit of relatively high weight is surprising, and hence is also remembered (though a very complex unit may be too difficult to remember well). This agrees with Meyer’s view [56] that anticipation plays a key role. The relation of these hypotheses to emotion is considered in Section 6.

The determination of saliency is of course an empirical issue, and will differ not only from one culture to another, but also from one style to another, and even from one listener to another. Given adequate saliency data, appropriate components and weights can be determined for a model, which could then be run and compared with the judgements of real listeners (it is of course necessary to carefully design experiments for this purpose). The flexibility to choose components and weights is an important feature of this approach, and distinguishes it from doomed modernist attempts to define universal aesthetic measures, such as that of George Birkhoff [8]. Just to be clear, the description of components is done using semiotic spaces, and the composition of a musical piece from its components is a blend of those components, in exactly the sense of structural blending that is described in Section 4.2. The cross-space mappings that result from blending are intra-musical metaphors, especially those that arise from transformations.

The Anticipatory Model captures aspects of Husserl’s phenomenology of time. For example, it has versions of both retention and protention, and the right kind of relationship between them. It also implies Husserl’s pithy observation that temporal objects (i.e., salient events or qualia) are characterized by both duration and unity. Since it is not useful to anticipate details very far into the future, because the number of choices grows very quickly, an implementation of protention, whether natural or artificial, needs a structure to accommodate multiple, relatively short projections, based on what is now being heard, with weights that increase with elapsed time; this is a good candidate for implementation by a neural net of competing Hebbian cell assemblies, in both the human and algorithmic instantiations, as well as robots (as in [71]), and it also avoids reliance on old style AI representation and planning.

The mathematical properties of the hierarchical complexity measure are very pleasing: it satisfies all the major equations and inequations of the classical Shannon information theory<sup>22</sup>, even though it greatly generalizes that notion, as well as generalizing Kolmogorov style complexity notions. This generalization justifies using such terms as “entropy” or (under a different metaphor) “temperature” for  $H(S)$ , although the term “density” might be preferred by musicologists. Moreover, colimits in the sense of category theory [55] are involved in the composition of hierarchical systems [31], in the blending of conceptual spaces, and more generally, of semiotic systems [34].

The **conditional complexity function**, denoted  $H(S'|S)$ , intuitively measures the *additional* effort needed to understand  $S'$  given that  $S$  is already understood [31], or the *novelty* of  $S'$  relative to  $S$ ; it provides a model for the cognitive distance<sup>23</sup> from  $S$  to  $S'$ . In the case of music where  $S', S$  are temporal series,  $S$  might be an initial segment of  $S'$ . Protention attempts to predict an extension  $S'$  of  $S$  out to some limit  $\delta$ . Given a piece using some fixed elements, transformations,

---

<sup>22</sup>However, these relations will only be approximately satisfied for the Anticipatory Model.

<sup>23</sup>Strictly speaking, it provides a quasi-pseudo-metric, since  $H(S'|S)$  is not symmetric, i.e., sometimes  $H(S'|S) \neq H(S|S')$ , and in general it also fails to satisfy the property  $H(S'|S) = 0$  iff  $S' = S$ , although it does satisfy this when restricted to series of the same length, in which case it gives a quasi-metric, which can then be symmetrized to yield a proper metric,  $d(S', S) = \max\{H(S'|S), H(S|S')\}$ .

and weights, a **complexity profile** is determined for the piece,  $\varphi(S, t) = H(S^t | S^{t-\delta})$ , where  $S^t$  denotes an initial segment of  $S$  up to time  $t$ , and  $\delta$  is a buffer size; intuitively,  $\varphi$  is intended to measure the cognitive effort required to understand the structure of the piece at each moment  $t$ . Sharp variations in  $\varphi(S, t)$  will correspond to the boundaries of important structural units. For example, in an AABA form, there will be a sharp drop in complexity at the boundary between the first and second A unit, and a sharp rise at the boundary between the second A and the B unit. Of course, the exact values will depend on what material is encoded in the description  $S$  of the piece, for example, whether micro-tonal and micro-timbral inflections are included.

Among many other concepts that can be defined within this theory, a measure of how much one structure  $S'$  resembles another  $S$ , given by  $r(S', S) = H(S'|S)/H(S)$ , seems relevant for thinking about improvisation, where  $S', S$  are segments over the same temporal interval, produced by two different improvisers, since  $r$  measures the extent to which one musician is following, or is being influenced by, the other. This **resemblance measure** is *not* symmetric, and hence is closer to our intuition for this application than the usual statistical concept of correlation, which is symmetrical. Two other important concepts generalizing classical Shannon theory, are the **mutual information** of  $S'$  and  $S$ , defined by  $I(S', S) = H(S') - H(S'|S)$ , and the **joint information** of  $S'$  and  $S$ , denoted  $H(S', S)$ , defined to be the minimum complexity required to realize *both*  $S$  and  $S'$ . Mutual and joint information are both symmetric in their two arguments; see [31] for more detail, including proofs of these and other equalities and inequalities familiar from the Shannon theory. Another interesting notion for exploration is the complexity difference between the anticipated and the actual continuation of a segment, since this measures the “surprise” of the actual continuation.

We intend to use the Anticipatory Algorithm in some computational experiments, to validate the assumptions behind it on real musical examples. We will begin with simple examples, and then building on that experience, gradually move to more and more complex examples. We have already studied some simple melodies (e.g., nursery rhymes) in [31], and will work our way up, with (for example) Charlie Parker solos as an intermediate step, towards contemporary group improvised music, some qualitative aspects of which are explored in work with David Borgo [10], as briefly discussed at the end of Section 5. It would also be interesting to analyze some of the examples used by Meyer [56]. This experimental programme should raise interesting challenges for the theory, and thus stimulate its further development. Some of those challenges will undoubtedly involve qualia.

Our complexity based approach can be extended to more directly address dynamic aspects of music, by viewing the space of pieces as a dynamical system<sup>24</sup>, with conditional complexity  $H(S'|S)$  providing a notion of distance (see footnote 23) and with the complexity profile  $\varphi$  as a potential or energy function, so that notions like curvature, basin of attraction, and saddle point become meaningful. Here are some examples of dynamical systems concepts that correspond to musical concepts: a basin of attraction is a region of low energy; a direction of motion (given by a tangent vector to a piece) is similar to the notion of “gesture” as used in musicology; complexity is similar to what is often called “density” in music; and phase transitions correspond to significant changes in the texture of segments. Note that the Anticipatory Model should be used, since the complexity measures are used to extrapolate along a given path, and that one can apply the same ideas to projections of the state space onto some lower dimensional space; with an appropriate choice of coordinates, this could give an interesting phase space for some class of pieces.

---

<sup>24</sup>This refers to a class of mathematical models of complex systems having origins mainly in physics, but with applications to many areas, including multimedia [4]. We will use terminology from this area; e.g., see [3] for a precise but relatively readable exposition. Due to their complexity, further details of this approach are deferred to a future paper; moreover, realizing its full potential will require non-trivial additions, such as extending the hierarchical complexity theory to continuous systems.

## 5 The Social Life of Music

Many writers have addressed the social aspects of music, among whom one might particularly mention Theodor Adorno [1]; of course, music is irreducibly social, and there is little reason to belabor this point here. However, it seems worth raising two issues which further illuminate some previous discussion in this paper. The first is improvisation, which draws on an essay by an improviser (Ryoko Goguen) on a particular performance (at which this author was also present), given in Appendix A. The second issue is the evolution of musical notation, with some emphasis on the values that this reveals.

Improvisation challenges notions of music as pre-existing ideal form, as score, or as a set of performances. Every improvisation is unique, and is uniquely tied to the particular environment in which it occurred. Thus, for the improvisation discussed in Appendix A, the location in Vienna, the start of the Iraq invasion, the atmosphere of a university seminar, the audience of advanced music students, and the sophisticated hosts, all played important roles in what was performed, how it was performed, and how it was received. Even exactly the same acoustic energy in a different environment would have been a different event, e.g., consider a cafe, a night club, and a street corner; and no doubt, each such change of environment would also change the performance. When a similar seminar was given at a different university in Vienna a few days later, already the feeling of the war was different, the audience was different (there were many undergraduate arts majors), the room was different (crowded), and so of course, the performance was different<sup>25</sup>.

This has profound implications for qualia: the qualitative feeling of music (or of anything else) is not just cognitive, it is also social. As anyone sufficiently experienced with improvisation knows, qualia arise in interaction with other musicians, the audience, the room (its particular acoustics, e.g., its resonant frequencies), the nature of the invitation to perform, the history of other performances in similar circumstances, and more, much more. As Ryoko wrote, “improvisation is interconnection.” In the technical language of Section 2.2, qualia are situated, and that situatedness is not just cognitive, or even embodied, it is enacted: Qualia are enacted in particular performances; they cannot be separated from the dynamics of those performances, and from everything that contributes to those dynamics. Following Vygotsky [77], we would especially mention material mediation (e.g., instruments, scores, gestures by musicians, and sound itself), and cultural, historical development, including the stylistic expectations residing in long term memory, which will of course change as listening experience accumulates.

Turning to the second topic, the notation of Western classical music has an interesting evolution, with a clear progression towards exercising more control over performers. Early music notation did not even specify duration, let alone amplitude, whereas some recent notation attempts to specify various aspects of intonation, although this can be very difficult, as the subtle inflections of jazz musicians such as Miles Davis make especially (and very beautifully) clear. Also, recall our earlier mention of Nono’s fermata.

Notation is never value neutral: it involves choices of what to notate, that is, of what is important; even the choice *to* notate involves the value of exercising certain forms of control over performers. The situation is similar for transcription, and even for analysis, because an analyst must make value-laden choices of what to analyze, how to analyze it, and how to report the results, among other things. For example, spectral analysis involves the value choice of leaving out performer interactions. However, this does not mean it is necessary to become hopelessly mired in a swamp of relativity. One can always speak in one’s own voice, from one’s own values, and one

---

<sup>25</sup>A reader who has pondered Section 4.3 might wonder how audience, room acoustics, etc. could enter into the Anticipatory Model to affect its “listening.” The answer is rather simple: an audience can be transcribed as if it were one of the musicians, and related to other time-varying parameters by appropriate components and weights.

can also compare various forms of analysis, and use that to help infer their values. One approach is to expose some details that are ignored by some form of analyses, but that are important for “alternative” forms of music (note the highly value-laden use of the term “alternative”!). For example, one might use closely timed sequences of spectral analyses instead of transcription as a basis for examining some fine details of pitch, intonation, etc. Of course any discussion of such data is itself a form of transcription, which again necessarily involves value choices.

Similar considerations apply to structure, which also encodes values, since deciding what counts as structure is a value-laden choice. For example, any preference between symmetry and asymmetry is value-laden, and even what counts as symmetry is culturally determined. The historical trend towards ever more precise notation, and the concomitant downgrading of improvisation within the classical tradition, attest to an ongoing devaluation of qualitative aspects of experience, in favor of more quantitative aspects, and this correlates with some of the philosophical movements discussed earlier in this paper. Qualia represent a kind of rebellion against all this, manifesting not only in philosophy, but also as a kind of social movement, e.g., [57]. But despite all this, it should certainly not be thought that qualia are incompatible with technical progress or with mathematical models; as suggested in Sections 4.2 and 4.3, such techniques can be used to model qualia, and to explore their properties in greater depth.

An interesting hypothesis is that good jazz improvisation involves forms of large grain structure that differ from those of classical music, and if so, we can hope for our minimum complexity algorithm to make these structures more explicit; there are also some interesting small grain structures to be explored. Of course, the very idea of looking for such structures involves values. A related idea is to identify patterns of interaction among improvising musicians. A research project on this could make dual use of video recordings and the structural resemblance measure  $r$  introduced in Section 4.3, to try to determine social and musical contexts where one improviser imitates, supports, or challenges another.

Application of the dynamical systems approach to free improvisation are being explored in collaboration with David Borgo [10]. A preliminary observation is that improvisers do not behave like particles in physics: they rarely linger in low energy basins of attraction, and often try to avoid even getting near them; they sometimes “defy gravity” (or “defy entropy,” depending on the choice of metaphor) and jump to a higher complexity; and they sometimes to surf an “edge of chaos” for a bit, before veering off in a new direction; see also [9]. No doubt the implicit drama and sense of freedom involved in this is part of the attraction of the best improvised music. It makes intuitive sense that the dynamic development of music should involve anticipation, time consciousness, qualia, and complexity in fundamental ways, and one might go so far as to suggest that theories that do not address these issues cannot be adequate.

## 6 Conclusions

This paper has attempted to explore the qualitative aspects of experience using music as data, and to place this exploration in the context of some relevant philosophical, cognitive scientific, and mathematical theories. Our observations have supported certain theories and challenged others. Among those supported are Husserl’s phenomenology of time, Vygotsky’s cultural-historical approach, and Meyer’s anticipatory approach, while Chalmers’ dualism, Brentano’s thesis on intentionality, qualia realism, qualia atomism, Hume’s pointilist time, and classical cognitivism have been disconfirmed at least in part. All these positions were discussed in the body of the paper, often briefly, and sometimes without the technical terminology. “Qualia realism” asserts that qualia have real existence (or “ontological status”), and “qualia atomism” asserts that qualia are indecomposable.

Among the new ideas in this paper, one might wish to highlight the discussion of context in music, the definition of qualia that avoids dualism, the application of structural blending to music, the Anticipatory Model of music understanding, the Qualia and S Hypotheses about the structure of qualia, the emphasis on social and value issues in music, and the research projects suggested by the resemblance measure and the complexity-based dynamical systems model.

Cognitivism was also mentioned several times but not discussed in detail. Like many other theories, cognitivism arose as a rebellion against the overly restrictive worldview of some prior theory, in this case, behaviorism, which tried to study behavior without invoking mind. Cognitivism in the broad sense of taking mind seriously, is admirable, but in fact, most cognitivist research takes a much more narrow view, in which cognition is considered computation, so that body, emotion, and society are neglected, and the representation of knowledge emerges as a central problem. In its classic form, now called “good old fashioned AI” or “GOFAI,” knowledge is represented in symbolic logic, an approach which the logical positivists of the Vienna Circle would presumably have endorsed. The conspicuous failure of this approach, e.g., in the Japanese Fifth Generation project, has inspired a number of biologically motivated refinements, such as neural nets and so-called artificial life, which do not, however, abandon the computational model, nor do they solve the problem of representation, which can be more precisely formulated as the *symbol grounding problem*, posed, but not solved, by Stevan Harnad [46]: the issue is how the symbolic representations used in a computational model can come to refer to the real world.

While the information processing models of cognitivism appear adequate for many aspects of low level perception, their exclusion (or cursory treatment) of embodiment, emotion, and society render them unsatisfactory as a theory of what it means to be human [24, 76]. In particular, embodiment, emotion, and society are certainly important parts of how real humans can be living solutions to the symbol grounding problem. The pervasive influence of cognitivism is presumably one reason why qualia in general, and emotion in particular, have been so neglected by traditional philosophy of mind, AI, linguistics, and so on. We may hope that this is now beginning to change.

Our observations on musical qualia undermine (computational) cognitivism. Descartes’ *Discourse on Method* [21] and nearly all subsequent methodological guidelines for scientific research, suggest that elements of any reasonable kind should combine in ways not too dissimilar from how atoms combine to form molecules. But we have seen that musical qualia behave very differently from that ideal, so much so that it is difficult to imagine rules for a “calculus of qualia” that could adequately describe how qualia combine. This is in part because qualia are temporal entities arising through the operation of a complex, heterogeneous system<sup>26</sup> that is largely unconscious, including modules that can be roughly described as buffers, memories, pattern recognizers, anticipators, and emotional feedback<sup>27</sup>. Although such descriptions using an information processing metaphor can certainly be helpful, the underlying complexity makes it impossible to do realistic simulations, and unreasonable to expect simple laws to hold for qualia. Indeed, reference to cultural norms, historical conditions and precedents, private emotional states, public reactions, and so on, are usually much more helpful in understanding music than speculations based on neural architecture. It is typical of emergent phenomena, i.e., of levels of complexity that can only be seen at a larger grain of description, that they do not fully reduce to lower levels; in this sense, they are *emergent*. For example, laws about valence in chemistry are difficult to reduce to the level of quantum mechanics. Because of such complexities, it seems worthwhile to explore the simpler, higher level models of Section 4. However, this does not mean that valuable insights cannot be had from models at the

---

<sup>26</sup>This is typical of “systems” produced by evolution, through the gradual accretion of new features on top of older structures; it also suggests that the music of a highly developed alien species would almost certainly be mutually unintelligible (and even ugly) to humans.

<sup>27</sup>Some interesting suggestions about the neural bases of qualia are given in [62].

neural level, as illustrated by the possible cell assembly model for protention in Section 4.3. On the other hand, the dynamical systems theory of Section 4.3 models qualia, emotion, etc. at a more abstract, emergent, global level.

Anticipation has been well explored in some areas of cognitive neuroscience, including efference copy in perception, and Daniel Dennett has even written that “all brains are, in essence, anticipation machines” [18]. However, cognitive science has paid little attention to the kind of abstract anticipation involved in music, let alone to its connection with emotion, and this paper also has so far said relatively little about emotion, which is after all the essence of qualia. Meyer [56] argues for a “law of affect,” that emotion in music “is evoked when a tendency to respond is inhibited,” or in a less behaviorist language, when an expectation is not met. However, this approach seems too simple, as is suggested by the numerous qualifications that Meyer is forced to invoke in order to apply this law to particular examples. We believe that the Qualia Hypothesis given in Section 4.3 is more precise, covers at least as many examples, and is more amenable to empirical test, though of course much work is needed to validate these beliefs. It seems possible that some such approach could provide a foundation for a future musicology that is both more precise and more adequate for contemporary forms of expression.

There is a clear evolutionary explanation for the role of emotion in helping organisms predict events in their environment: To promote survival, moderate curiosity and discovery should be rewarded, while uncertainty should arouse interest, but should become unpleasant when sufficiently strong. Let us call this **Hypothesis E**. Combined with the Qualia Hypothesis, it provides a rough guide to the kinds of emotion in music, and the circumstances of their production: emotion arises from relations in the “now” among retention, protention, and perception; the right balance of novelty and predictability will maintain interest and pleasure, while too little novelty will cause loss of interest, and too much will become unpleasant<sup>28</sup>.

The hypotheses in this paper have other implications, but to avoid diluting its focus, we mention only the following. They explain why music exhibits a continual development of new ways to create and manipulate listener expectations, as older ways become too familiar; indeed, it is possible to view the whole history of music in this light. They also explain the pleasure that musicians (and many other professionals) take in learning new skills and further developing old skills. Another implication for learning is that test-oriented teaching can be anti-productive, because it can inhibit emotional response, prevent exploration, and dull curiosity.

Consciousness has so far been mentioned mainly in relation to time and qualia. Section 2.2 and Section 3 argued that enactment not only constructs objects-in-time, but also their contexts, and even time itself, through processes that involve retention, and protention or anticipation. Section 4.3 argued that qualia are determined by variations of saliency, or in the model, of weight, as an intrinsic part of these processes, which it therefore seems reasonable to call **semiosis**. Both Husserl’s time construction and our qualia construction yield conscious events, and the former can be seen as implied by the latter, although it should be noted that much of what goes on is unconscious, including much of anticipation and saliency determination.

This suggests that *all* consciousness arises through sense-making processes involving anticipation, which produce qualia as sufficiently salient chunks. Let us call this the **C Hypothesis**; it provides a theory for the origin and structure of consciousness. If correct, it would help to explain why consciousness continues to seem so mysterious: it is because we have been looking at it through inappropriate, distorting lenses; for example, attempting to view qualia as objective facts, or to assign them some new ontological status, instead of seeing segmentation by saliency as an inevitable feature of our processes of enacted perception. It also implies that much research on

---

<sup>28</sup>But it should be noted that much contemporary music is exploring gradations of this unpleasantness.

consciousness has been misfocused on peripheral issues, thereby producing misleading results. In addition, it suggests that music is an ideal experimental medium for the study of consciousness, since it allows us see fine details in the production of qualia, is subject to experimental manipulation in a variety of ways, already has a rich literature and discipline of training, and is amenable to model construction at several levels<sup>29</sup>. The Main, Qualia, E and C Hypotheses help understand the structure of consciousness: its segments are qualia, hierarchically organized by their saliency, with emotional tone determined by their resonance with protention. It will be surprising if the future does not see a great deal more research on time, emotion and consciousness using music as data, as well as research on music using concepts from cognitive science and consciousness studies.

Many writers have advanced ideas that are consistent with and/or similar to the C Hypothesis; we now briefly survey just a few of these. The role of time in consciousness was stressed by William James, who described the “specious present” as having “a certain breadth of its own on which we sit perched, and from which we look in two directions into time” [50]; James thought this was about 12 seconds, in close agreement with [66]. James also wrote that the stream of consciousness is “like a bird’s life, made of an alternation of flights and perchings,” a very fine description of the chunking of consciousness, and he considered the role of anticipation, using the term “tendency.” Martin Heidegger’s ready-to-hand [47] is also related. However, the combination of consciousness and anticipation has received relatively little scientific attention until recently. Ralph Ellis and Natika Newton [25] advances a hypothesis related to ours, but do not consider time, qualia, or their structures. Moreover, Ellis and Newton, and many others, including Rodney Cotterill [13], Rodolfo Llinás, and Gerald Edelman, limit their consideration of anticipation to sensory-motor systems, which excludes the case of music.

An insightful paper of Newton [60] comes even closer, suggesting a naturalistic explanation for the emergence of consciousness, in which brain processes provide a physical basis for the mysterious and ineffable character of qualia, without claiming that they actually exist. Newton suggests that consciousness emerges to handle the novelty that results from blending incompatible components. However, her notion of blending is binding, whereas our notion is a precise mathematical formalization [34] of a structural generalization of the cognitive linguistic notion of blending<sup>30</sup>. Moreover, we do not agree with [60] that the blended items must be incompatible, claiming instead that more complex blends produce more conscious qualia. For us, ineffability arises because qualia have hierarchically structured parts with variable saliency, while temporal thickness results from the concurrent operation of perception, retention, and protention, whereas Newton attributes both ineffability and temporal thickness to the blending of incompatibles.

Although enthusiastic about anticipation, Dennett does little with it; moreover, his heterophenomenology and his arguments against qualia distance him from consciousness almost as far as the behaviorists put themselves. Gibson’s ecological perception is one early theory using anticipation in a serious way [29, 30]. There is also work on anticipation in the semiotic tradition by Mihai Nadin [58]. In my opinion, the insights of Husserl and James are the most suggestive. Husserl achieves a stunning depth of analysis of the phenomenology of time, while James exhibits a remarkable breadth and clarity of vision, in both cases without benefit of the enormous recent advances in cognitive science. But it should also be noted that many of the same insights can be found centuries earlier in the Buddhist meditation literature. Our approach differs from all these in its grounding in the Main Hypothesis, with its hierarchical networks of weighted, shared, and transformed components, and in the precision that that model can confer on predictions that are made using the other hypotheses.

---

<sup>29</sup>Note the consistency of these points with the fourth person method of Section 2.2.

<sup>30</sup>But under the perhaps presumptive claim that structural blending explicates binding, our notion can be seen as an explication of hers.

I wish to close this essay by extending a discussion about art and the sacred begun in [33]. The present paper has argued that consciousness, qualia, and the phenomenal world of experience are produced by complex and largely unconscious processes, but it has not pointed out that this undermines belief in the existence of an independent unified “self.” Our hypotheses do not imply that people do not exist, or that the appearance of a unified independent self does not exist, but rather that the unity and the independence have an illusory character. We are complex, dependent beings, continually recreated as “conscious” through salient events, determined as such in a complicated process that is cultural and historical, as well as cognitive. Similarly, our experience of the world is “manufactured” by this process, a complex and ever shifting mosaic of qualia of various textures and sizes; this does not imply that there is no “real world,” but rather that our access to it is achieved in a very complex and incomplete way. Similar insights were expressed nearly two thousand years ago by the Indian philosopher-sage Nagarjuna [59], although of course using a different language, that of Mahayana Buddhism. More recent thinkers like Dogen [22] and Keiji Nishitani [61], among many others, have also expressed similar insights.

My meditation teacher, Chögyam Trungpa Rinpoche, said that the purpose of art is to show our non-existence in the world. It seems to me that close attention to music, and to how we hear it, can give precisely this insight. The experience can be vivid, clear, and deeply moving. Indeed, it can be an experience of the “fourth moment” (*dus bzhi mnyam pa chen po* in Tibetan), an experience of time suspended, of being not past, present, or future, but a limitless space of great equanimity that unifies and transcends all three, and in which both self and world disappear. This space is the abode of the sacred.

## A “A” Improvisation Notes

The following gives a performer’s view of one particular improvisation. It is used as data in the body of this paper.

On 20 March 2003, the day the Iraq war began, I performed an improvisation at the University of Vienna (and again on 28 March, at the Performing Arts University of Vienna), in connection with lectures by my husband on some cognitive and philosophical aspects of music. Part of the lecture showed how the same “middle A” could be perceived totally differently, depending on its context.

I decided to use this same note, and some of those contexts, as a basis for improvisation, not forgetting Vienna, a city that experienced much pain as well as much great music in its history, and of course being very aware of the war, and thinking how everything is connected, the performer, the audience, the city, the war, the sadness and pain, the whole of history.

Middle A is a beginning sound, and on a good piano (such as the Bosendorfer grand they gave me for practice) you can hear its many colors hanging in space, carrying the pain of numerous lives, from Mozart to Webern, and now poor Turkish immigrant workers, and Iraqi civilians. But you can also hear hope and love.

Low A has an angry sound, an aggressive wish to control, a cynical view — but still with pain inside of it, and the possibility of understanding.

High A is like a scream, a loved one lost in the war, frustration, need; or a whisper, even a caress.

Every note has its overtones, and they are all connected, as are all of us, middle, high, low, hope, fear, pain, and love.

Improvisation is interconnection.

Ryoko Amadee Goguen  
3 June 2003

## References

- [1] Theodor Adorno. *Essays on Music*. California, 2002. Translated by Richard Leppert.
- [2] Edward Aldwell and Carl Schachter. *Harmony and Voice Leading*. Harcourt, 1989. Second edition.
- [3] Kathleen Alligood, Tim Sauer, and James Yorke. *Chaos: An Introduction to Dynamical Systems*. Springer, 1996. Textbooks in Mathematical Sciences.
- [4] Peter Bogh Andersen. Multimedia phase spaces. *Multimedia Tools and Applications*, 6:207–237, 1998.
- [5] Gerard Assayag, Hans Feichtinger, and José-Francisco Rodrigues. *Mathematics and Music: A Diderot Mathematical Forum*. Springer, 2002.
- [6] Roland Barthes. *Elements of Semiology*. Hill and Wang, 1968. Trans. Annette Lavers and Colin Smith.
- [7] Jon Barwise and John Perry. *Situations and Attitudes*. MIT (Bradford), 1983.
- [8] George David Birkhoff. *Aesthetic Measure*. Harvard, 1933.
- [9] David Borgo. The orderly disorder of free improvisation. *Pacific Review of Ethnomusicology*, 10, 2002.
- [10] David Borgo and Joseph Goguen. Sync or swarm: Group dynamics in musical free improvisation. In Richard Parncutt, Annekatrin Kessler, and Frank Zimmer, editors, *Proceedings, Conference on Interdisciplinary Musicology*, pages 52–52. 2004. Held Graz, Austria, 15–18 April 2004.
- [11] David Chalmers. *The Conscious Mind: In search of a fundamental theory*. Oxford, 1996.
- [12] Paul Churchland. *The Engine of Reason, the Seat of the Soul*. MIT, 1996.
- [13] Rodney Cotterill. Evolution, cognition and consciousness. *Journal of Consciousness Studies*, 8(2):3–17, 2001.
- [14] Tim Crane. *The Mechanical Mind*. Routledge, 2003. Second edition.
- [15] Francis Crick. *The Astonishing Hypothesis: The scientific search for the soul*. Scribner & Sons, 1994.
- [16] Antonio Damasio. *Descartes' Error: Emotion, Reason and the Human Brain*. Avon, 1994.
- [17] Antonio Damasio. *The Feeling of What Happens: Body, Emotion and the Making of Consciousness*. Harcourt, 1999.
- [18] Daniel Dennett. *Conscious Explained*. Penguin, 1991.
- [19] Daniel Dennett. Quining qualia. In Edoardo Bisiach and Anthony Marcel, editors, *Consciousness in Contemporary Science*, pages 42–77. Oxford, 1996. Reprint of 1988 edition.
- [20] Daniel Dennett. Who's on first? Heterophenomenology explained. *Journal of Consciousness Studies*, 10(9–10):19–30, 2003.
- [21] René Descartes. *Discourse on Method*. Prentice Hall, 1956. Original from 1637.
- [22] Dogen. *Shobogenzo: Zen Essays by Dogen*. University of Hawaii, 1992.

- [23] Paul Dourish. What we talk about when we talk about context. *Personal and Ubiquitous Computing*, 2003. Online edition.
- [24] Hubert Dreyfus. *What Computers Still Can't Do*. MIT, 1992.
- [25] Ralph Ellis and Natika Newton. Three paradoxes of phenomenal consciousness: Bridging the explanatory gap. *Journal of Consciousness Studies*, 5(4):419–442, 1998.
- [26] Gilles Fauconnier and Mark Turner. Conceptual integration networks. *Cognitive Science*, 22(2):133–187, 1998.
- [27] Gilles Fauconnier and Mark Turner. *The Way We Think*. Basic, 2002.
- [28] Harold Garfinkel. *Studies in Ethnomethodology*. Prentice-Hall, 1967.
- [29] James Gibson. The theory of affordances. In Robert Shaw and John Bransford, editors, *Perceiving, Acting and Knowing: Toward an Ecological Psychology*. Erlbaum, 1977.
- [30] James Gibson. *An Ecological Approach to Visual Perception*. Houghton Mifflin, 1979.
- [31] Joseph Goguen. Complexity of hierarchically organized systems and the structure of musical experiences. *International Journal of General Systems*, 3(4):237–251, 1977.
- [32] Joseph Goguen. Semiotic morphisms, 1996. Available on the web at [www.cs.ucsd.edu/users/goguen/papers/smm.html](http://www.cs.ucsd.edu/users/goguen/papers/smm.html). Early version in *Proc., Conf. Intelligent Systems: A Semiotic Perspective, Vol. II*, ed. J. Albus, A. Meystel and R. Quintero, Nat. Inst. Science & Technology, (Gaithersberg MD, 20–23 October 1996) pages 26–31. See also UCSD Dept. Computer Science & Eng. Report CS97–553, 1997.
- [33] Joseph Goguen. Introduction. *Journal of Consciousness Studies*, 6(6/7):5–14, 1999. Introduction to special issue, *Art and the Brain, Part 1*.
- [34] Joseph Goguen. An introduction to algebraic semiotics, with applications to user interface design. In Chrystopher Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 242–291. Springer, 1999. Lecture Notes in Artificial Intelligence, Volume 1562.
- [35] Joseph Goguen. What is art? *Journal of Consciousness Studies*, 7(8/9):7–15, 2000. Introduction to special issue, *Art and the Brain, Part 2*.
- [36] Joseph Goguen. Semiotic morphisms, representations, and blending for interface design. In *Proceedings, AMAST Workshop on Algebraic Methods in Language Processing*, pages 1–15. AMAST Press, 2003. Conference held in Verona, Italy, 25–27 August, 2003.
- [37] Joseph Goguen and Fox Harrell. Information visualization and semiotic morphisms. In Grant Malcolm, editor, *Visual Representations and Interpretations*. Elsevier, 2003. Proceedings of a workshop held in Liverpool, UK.
- [38] Joseph Goguen and Fox Harrell. Foundations for active multimedia narrative: Semiotic spaces and structural blending, 2004. To appear in **Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems**.
- [39] Joseph Goguen and Charlotte Linde. Optimal structures for multi-media instruction. Technical report, SRI International, 1984. To Office of Naval Research, Psychological Sciences Division.
- [40] Joseph Goguen and Grant Malcolm. *Algebraic Semantics of Imperative Programs*. MIT, 1996.
- [41] Joseph Goguen, Grigore Roşu, and Kai Lin. Conditional circular coinductive rewriting. In *Recent Trends in Algebraic Development Techniques, 16th International Workshop, WADT'02*. Springer, Lecture Notes in Computer Science, to appear 2003. Selected papers from a workshop held in Frauenchiemsee, Germany, 24–27 October 2002.

- [42] Joseph Goguen, James Thatcher, and Eric Wagner. An initial algebra approach to the specification, correctness and implementation of abstract data types. In Raymond Yeh, editor, *Current Trends in Programming Methodology, IV*, pages 80–149. Prentice-Hall, 1978.
- [43] Nelson Goodman. *Languages of Art: An Approach to a Theory of Symbols*. Bobbs-Merrill, 1968.
- [44] Joseph Grady, Todd Oakley, and Seanna Coulson. Blending and metaphor. In Raymond Gibbs and Gerard Steen, editors, *Metaphor in Cognitive Linguistics*. Benjamins, 1999.
- [45] Eduard Hanslick. *On the Musically Beautiful*. Hackett, 1986. edited by Geoffrey Payzant.
- [46] Stevan Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.
- [47] Martin Heidegger. *Being and Time*. Blackwell, 1962. Translated by John Macquarrie and Edward Robinson from *Sein und Zeit*, Niemeyer, 1927.
- [48] David Hume. *A Treatise of Human Nature*. Viking, 1986. Original from 1740.
- [49] Edmond Husserl. *Phenomenology of Internal Time-Consciousness*. Indiana, 1964.
- [50] William James. *Principles of Psychology*. Dover, 1950. Originally published 1890.
- [51] Peter Kivy. *The Corded Shell: Reflections on Musical Expression*. Princeton, 1980.
- [52] Andrei Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1:1–11, 1965.
- [53] George Lakoff. *Women, Fire and Other Dangerous Things: What categories reveal about the mind*. Chicago, 1987.
- [54] George Lakoff and Mark Johnson. *Metaphors We Live By*. Chicago, 1980.
- [55] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 1971.
- [56] Leonard Meyer. *Emotion and Meaning in Music*. Chicago, 1956.
- [57] Ken Mogi. Qualia manifesto. <http://www.qualia-manifesto.com>.
- [58] Mihai Nadin. *Mind: Anticipation and Chaos*. Belser, 1991.
- [59] Nagarjuna. *Mulamadhyamikakaraka*. Oxford, 1995. Translated by Jay Garfield.
- [60] Nitika Newton. Emergence and the uniqueness of consciousness. *Journal of Consciousness Studies*, 8(8–9):47–59, 2001.
- [61] Keiji Nishitani. *Religion and Nothingness*. University of California, 1982.
- [62] Georg Northoff. Qualia and ventral prefrontal cortical function: ‘neurophenomenological’ hypothesis. *Journal of Consciousness Studies*, 10(8):14–48, 2003.
- [63] Charles Saunders Peirce. *Collected Papers*. Harvard, 1965. In 6 volumes; see especially Volume 2: Elements of Logic.
- [64] Plato. *The Republic*. Everyman, 1993.
- [65] Vilayaner S. Ramachandran and William Hirstein. The science of art: a neurological theory of aesthetic experience. *Journal of Consciousness Studies*, 6(6/7):15–51, 1999.
- [66] Mikko Sams, Riitta Hari, J. Rif, and J. Knuutila. The human auditory sensory memory trace persists about 10 sec: Neuromagnetic evidence. *Journal of Cognitive Neuroscience*, 5:363–370, 1993.
- [67] Ferdinand de Saussure. *Course in General Linguistics*. Duckworth, 1976. Translated by Roy Harris.
- [68] Claude Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois, 1964.

- [69] Lucy Suchman. *Plans and Situated Actions: The Problem of Human-machine Communication*. Cambridge, 1987.
- [70] David Sudnow. *Ways of the Hand*. Bantam, 1979. Reprinted by Harvard.
- [71] Jun Tani. The dynamical systems accounts for phenomenology of immanent time: An interpretation from a robotics synthetic study. *Journal of Consciousness Studies*, 11:Submitted for publication, 2004.
- [72] Richard Taylor. The use of science to investigate Jackson Pollock's drip paintings. *Journal of Consciousness Studies*, 7(8/9):137–150, 2000.
- [73] Mark Turner. *The Literary Mind*. Oxford, 1997.
- [74] Francisco Varela. Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies*, 3(4):330–349, 1996.
- [75] Francisco Varela. Present-time consciousness. *Journal of Consciousness Studies*, 6(2/3):111–140, 1999.
- [76] Francisco Varela, Evan Thompson, and Eleanor Rosch. *The Embodied Mind*. MIT, 1991.
- [77] Lev Vygotsky. *Mind in Society*. Harvard, 1985.
- [78] Alfred North Whitehead. *Process and Reality*. Free, 1985. Originally published 1929.
- [79] Gerhard Widmer. Discovering simple rules in complex data: a meta-learning algorithm and some surprising musical discoveries. *Artificial Intelligence*, 146(2):129–148, 2003.
- [80] Ludwig Wittgenstein. *Tractatus Logico-Philosophicus*. Routledge and Kegan Paul, 1922. English translation by D.F. Pears and B.F. McGuinness, with an Introduction by Bertrand Russell; original German edition in *Annalen der Naturphilosophie*, 1921.
- [81] Lawrence Zbikowski. *Conceptualizing Music*. Oxford, 2002.
- [82] Semir Zeki. Art and the brain. *Journal of Consciousness Studies*, 6(6/7):76–51, 1999.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Some Philosophy of Music</b>	<b>2</b>
2.1	Some Traditional Western Views . . . . .	2
2.2	Context, Situatedness, and Embodiment . . . . .	3
<b>3</b>	<b>Musical Qualia</b>	<b>5</b>
3.1	Time Consciousness . . . . .	7
<b>4</b>	<b>New Models for Understanding Music</b>	<b>9</b>
4.1	Metaphor and Blending . . . . .	9
4.2	Semiotic Spaces and Structural Blending . . . . .	11
4.3	Structural Complexity, Protention and Qualia . . . . .	13
<b>5</b>	<b>The Social Life of Music</b>	<b>17</b>
<b>6</b>	<b>Conclusions</b>	<b>18</b>
<b>A</b>	<b>“A” Improvisation Notes</b>	<b>22</b>