

ASPECTS OF MEANING CONSTRUCTION IN MUSIC: TOWARD A COGNITIVE GRAMMAR OF MUSIC

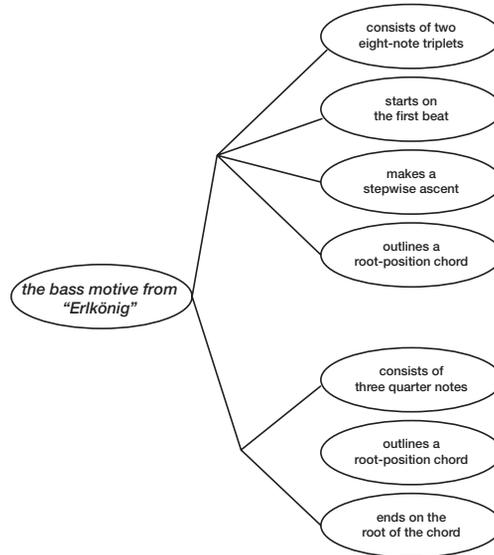
LAWRENCE M. ZBIKOWSKI

University of Chicago, Department of Music¹

Abstract: Operating under the assumption that human meaning is a product of human cognitive processes, this essay explores the source and nature of the concepts basic to our understanding of music. The essay shows how recent research on categorization, cross-domain mapping (including conceptual blending), and conceptual models can be used to account for the origin and structure of musical concepts. Using these ideas, the essay explains why musical works can seem fixed even when they are endlessly mutable

Keywords: music; categorization; cross-domain mapping; metaphor; conceptual blending; conceptual models; mental models; frames; musical ontology; blues

The scene is one familiar to many of us: the accompanist poised at the piano, the singer, framed by the lights, momentarily frozen in crystalline silence. And then, we hear the introduction and opening stanza of Franz Schubert's early song, "Erlkönig" (see Ex. 1). The world Schubert opens up in little more than a minute of music is a powerful one, and unique: there is nothing else quite like it in our daily experience. The mysterious power of music to work such magic has led people living in different moments of history representing different cultural perspectives to propose that it is something not of human origin. Music is over here, humans are over there.



Example 1: Score for introduction and first stanza of Franz Schubert's "Erlkönig," Op.-1

The fact remains, however, that we have little evidence of sustained musical activity from any source other than humans. And so the question becomes, how does this unique kind of experience fit in with the rest of our lives? How do we conceive music, and how do we conceive of it? How does music come to have the meaning it has for us?

In what follows I would like to approach these questions from the perspective of recent research in cognitive science. This research is remarkable for one important reason: it operates on the assumption that the answers are to be found in the way humans construct their understanding of the world. This is a departure from earlier theories of knowledge, which assumed that the answers to questions about meaning were to be found in relationships between symbolic structures and objective reality. As Mark Turner has noted, within such theories,

Meaning is conceived ... as essentially anchored in states of affairs in objective reality, with the consequence that the meaning of an utterance must be the reality to which it refers. This leaves the human person out of the loop altogether: A semantic express train shoots straight from ... symbols to an objective reality without passing through the human brain, let alone stopping in the human brain, let alone taking its entire journey there.²

In contrast, the perspective that I want to present—a perspective that draws on the work of researchers like Turner, as well as the linguists George Lakoff, Gilles Fauconnier, and Ron Langacker, the philosopher Mark Johnson, the cognitive psychologist Larry Barsalou, and the neurologist Antonio Damasio—assumes that human meaning

is completely and irrevocably tied to human cognitive processes. This perspective fundamentally reshapes how we ask questions of the sort I proposed above, for it requires that we dispense with any concern about the way symbolic structures might represent objective reality. Instead, we need to turn our attention to the mental images or neural patterns that arise during the perceptual-motor processing of a particular event on the part of some particular human. I should emphasize that there is no expectation that such images or patterns necessarily have any fidelity to the objects they represent: they are as much creations of the brain as they are products of the external reality that prompts their creation.³

In my own inquiries into how human cognition shapes both music and our understanding of music, I have found it useful to focus on the level of the concept, understood as the most basic cognitive structure regularly accessible to thought. In the first three sections of this paper I want to outline the approach to conceptualizing music I have developed, and give some indication as to how it might apply to other expressive media.⁴ In the last portion I take this perspective and show how it can contribute to the development of a cognitive grammar of music, and exercise that reveals still further what cognitive science can tell us about the arts.

Categorization and Conceptualization

Concepts and Categories

Some have argued, as have Ray Jackendoff and Mark DeBellis, that music is non-conceptual.⁵ As a musician, this troubles me, since I like to think that my music making, even at its most impassioned and improvisational, has something to do with concepts.⁶ I do not think these concepts can always be reduced to words, but that is one of the reasons we have music: to express things that cannot be expressed with words alone.

As an example of a musical concept, let us consider the bass motive that first appears in measures-2 through 3 of Schubert's song (see Ex. 1). This motive, or some version of it, also appears in measures 4, 9, 11, 15, 17, 19, 24, 26, 28, 33, and 35. Clearly, the motive is important for understanding the song: it is compact and immediate; it presents each of the harmonies used to structure this passage; it provides a contrast to the steady triplet rhythm of the right hand; and it sets out a two-bar hypermeter that enacts a dialogue with the harmonic rhythm set up by cadential arrivals. These features of the motive help guide our thought about the song in two important ways: first, they provide a basis for remembering past events within the song, and second, they make it possible to link this particular concept with other concepts that appear in the opening measures—in particular, the repeated-note pattern of the right

hand, and the harmonic progression leading to the cadence in measure 8. There is, however, one significant complication with regarding the bass motive as a concept, and that is the difficulty of defining just what constitutes this concept. There are at least three different versions of the motive in this passage alone, and Schubert seems to vary the features of the motive quite deliberately. If the bass motive is a concept, it certainly seems different from ostensibly stable concepts like cup or spoon. (I should also note that in actuality the bass motive is a complex concept—that is, one that embraces other concepts. This will become more clear in what follows.)

In thinking about musical concepts, I have found it useful to draw on recent research on processes of categorization, with particular emphasis on what are often called natural categories, so called because they result from humans' interactions with their natural environments. There are three benefits to this approach. First, categorization is among the most basic of cognitive processes, and categories appear to be allied with concepts. In fact, some researchers draw no fundamental distinction between categories and concepts.⁷ Second, because natural categories arise from humans' interactions with their environments, it is not necessary to assume prior structuring by language—indeed, significant aspects of the process of categorization appear to be independent of language.⁸ Third, natural categories, without exception, have a graded structure: membership in the category is not fixed, but is graded through a dynamic process in which the attributes of potential category members are compared with the attributes most typically found within the category.⁹

Because graded structure will be important to my discussion of musical concepts, let me explore it in a bit more detail, taking as my example the category bird. Experimental rankings show that subjects view robins and sparrows as the best examples of birds, with owls and eagles lower down in the rankings and ostriches, emus, and penguins among the worst examples.¹⁰ All of these creatures are considered members of the category bird, but some better represent the category than others. Category structure is consequently graded according to typicality: category members range from the most typical to the least typical, with the former securely inside the bounds of the category (robins and sparrows) and the latter in danger of being excluded from the category (emus and penguins).

Explaining the graded structure of natural categories has proved a challenging task for researchers into categorization; perhaps the best way to account for this structure is to understand natural categories as organized around conceptual models. Conceptual models are relatively basic cognitive structures that act as guides for reasoning and inference. I want to consider conceptual models in greater detail below; for now it will be enough to say that conceptual models consist of concepts in specified relationships, which pertain to a specific domain of knowledge. In a natural category, the

conceptual model provides a simplified representation of category structure that incorporates knowledge about what values are most typical for a select group of attributes for the given category. This representation is then used to evaluate potential category members: if our conceptual model for the category bird is framed around the features of robins and wrens and we encounter a penguin, we will most likely conclude that, whatever else a penguin is, it isn't a very good example of a bird. Needless to say, the conceptual model for any category will be informed by both culture and geography, for these will play a part in determining what is considered a typical member of the category.

Returning now to the bass motive from "Erlkönig," let us approach the twelve statements of the motive that appear in Example-1 as analogous to a natural category.¹¹ A provisional conceptual model for the category is given in Figure 1.12 The model details the attributes that are most typical of the category, which include: (1) an ascending six-note figure that comprises two triplets, starts on beat 1, makes a stepwise ascent, and outlines a root-position chord; (2) a descending figure that is stated in quarter notes, outlines a root-position chord, and ends on the root of the chord. This model gives us a basis for evaluating individual members of the category. For instance, the initial version of the motive accords completely with the model, and so will be judged typical of the category. The version of the motive that first appears in measure-15, on the other hand, is less typical of the category: although it conforms with the first main component of the conceptual model, it conforms rather less well with the second. Finally, the version of the motive that first appears in measure-24 falls somewhere in between these two: while it conforms with the first main component, it diverges from one aspect of the second main component—the outlining of a root-position chord..

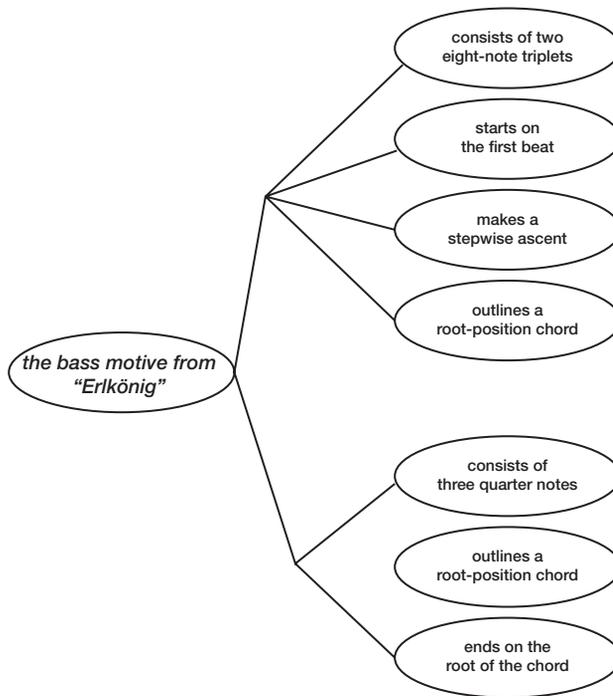


Figure 1: Provisional conceptual model for the category bass motives from Schubert's "Erlkönig"

The conceptual model of Figure-1, then, provides an idealized representation of the category bass motives from Schubert's "Erlkönig". The model, 'taught' to us by Schubert with his four iterations of the motive in the opening twelve measures, is stored in memory and used not only for the evaluation of further statements of the motive but for later hearings of the composition. The model is shaped in part by our global expectations for this sort of music—for instance, that it will have themes and motives prominently stated near the opening of the work—but is also in rapport with individual category members, which either corroborate or challenge the model. Through his deployment of different versions of the motive, Schubert hints at a process of departure and return: after establishing the model for the category in the opening measures, he moves in measure 15 to the most atypical form of the motive, proceeds from there to a somewhat more typical form, and finally returns to the original version of the motive.¹³

Categories, Concepts, and Syntax

What emerges from this discussion is an intriguing view of musical concepts and musical syntax, a view not only important for what follows, but also for expressive media other than music. Let me first turn to the notion of a musical concept. I realize that to attempt any definition of a concept is to embroil myself in a long-standing debate within philosophy and cognitive science.¹⁴ However, I want to skirt this debate by concentrating on a definition of “concept” that reflects recent work in the mind sciences and the brain sciences, and that is adequate to music. While I think this definition can inform the larger debate on ‘concept,’ I shall leave the substance of the debate for another time.

The definition of “concept” I want to develop here is influenced by the work of Gerald Edelman, who has been interested in developing a biological approach to consciousness. Edelman’s definition of a concept is developed as a description of the capacities necessary for the control of complex interactions between an organism and its environment:

An animal capable of concepts is able to identify a particular thing or action and control its future behavior on the basis of that identification in a more or less general way. It must act as if it could make judgments based on recognition of category membership or integrate “particulars” into “universals.” This recognition rests not just on perceptual categorization (although a concept may have a highly sensory content) but, to some degree, must also be relational. It can connect one perceptual categorization to another even in the absence of the stimuli that triggered these categorizations.¹⁵

Thus to have concepts involves not only the process of categorization, but also recognizing relationships between categories.¹⁶ What is also important is that having concepts is a capacity that is not limited to humans, a point also made by the zoologist Donald Griffin.¹⁷ Concepts are thus not necessarily tied to language.

With this perspective in mind, let me propose that a musical concept has three characteristics. First, it is a product of the process of categorization. A musical category is quite literally where our conceptualization of music begins. Second, a musical concept is an essential part of the means through which we guide present and future actions. These actions thus constitute a sort of indirect evidence for a cognitive structure almost as ephemeral as music itself. Third, a musical concept can be related to other concepts, including concepts associated with bodily states (both physical and emotional), perceptual categories (including sound, which, after all, is not necessarily music), and linguistic constructs.

From the perspective provided by this definition, language—in the sense of natural language—is not required in order to have musical concepts. I have, of course, used

language to characterize musical categories as well as the conceptual models around which musical categories are organized. But we could imagine a listener—perhaps a particularly astute listener, but perhaps simply an attentive amateur—who develops a musical category without recourse to language. While language is still nearly indispensable for communicating features of musical concepts and for developing such concepts in richer contexts, it is not required.¹⁸ The notion of concepts independent of language is more than a little intriguing, and with consequences that extend beyond music and the general debate on “concept” noted above. Thus where some authors have excluded music from the conceptual realm (because, by their definition, concepts require language),¹⁹ music can join dance, the visual arts, and any of the other non-linguistic modes of human expression as a properly conceptual activity.

This perspective on musical understanding also offers a new way to think about musical syntax. As I noted, the way Schubert deploys the bass motive from “Erlkönig” hints at a process of departure and return. Musical syntax can thus be viewed as an emergent property of sequences of musical events. I want to develop this approach to musical syntax further in the last portion of this paper; for now, I only want to emphasize how important categorization is not only for making sense of the musical moment, but for making sense of musical discourse as a whole.

Summary

Categorization, a very general process through which we structure our experience of the world as a whole, is something we also use to structure our understanding of music. Indeed, we cannot think of music without thinking in terms of categories. The knowledge associated with such categories—knowledge which, I should hasten to point out, is really rather complex, and thus occupies a quite high level in cognitive structure—is where musical meaning starts. This is not to say we are conscious of musical categories, for most of the time they go unnoticed, as do the overwhelming number of cognitive categories. It is only when something calls into question our normal way of categorizing things—a musical work unfamiliar to us, or a novel theoretical construct—that we become aware of how pervasive and important categorization is.

As I have mentioned, I shall return to the role of categorization in meaning construction in music later in this paper, and before that I shall have more to say about conceptual models (which provide the framework for categorical judgments). Right now, however, I want to turn to the question of how musical categories—or as I would regard them, musical concepts—connect with other aspects of our experience. To answer this, I would like to turn to work on cross-domain mapping.

Musical Concepts and Cross-Domain Mapping

Cross-Domain Mapping and Metaphor

As you may know, it was an exploration of metaphor theory initiated by Mark Johnson and George Lakoff in 1980 that led to the proposal that cross-domain mapping was a process basic to human thought. Through this process, we structure our understanding of one domain, which is typically unfamiliar or abstract, in terms of another, which is most often familiar and concrete. A standard example is our characterization of musical pitch: we speak of pitches as “high” and “low,” organizing our understanding of the ephemeral and nearly intangible acoustic domain in terms of orientation in physical space. In accordance with this mapping, I characterized the first main component of the model outlined in Figure-1 as an “ascending six-note figure,” even though no literal ascent is involved—the pianist’s left hand simply moves to the right to play the appropriate notes. The process of cross-domain mapping provides one way of articulating our musical concepts by placing them in correspondence with structures from other domains.²⁰

Fundamental to the theory of metaphor that sprang from Lakoff and Johnson’s work is a distinction between conceptual metaphors and linguistic metaphors. A conceptual metaphor is a cognitive mapping between two different domains; a linguistic metaphor is an expression of such a mapping through language. Thus the conceptual metaphor PITCH RELATIONSHIPS ARE RELATIONSHIPS IN VERTICAL SPACE, which maps spatial orientations such as up-down onto the pitch continuum, is the source of the metaphor in linguistic expressions like “On the guitar, the B string is higher than the G string.” As can be seen from looking at a guitar as it is held in playing position, although the pitch of the B string is “higher” than that of the G string, it is actually closer to the floor—that is, “lower,” than the G. I must hasten to add that, while conceiving of pitch relationships in terms of “high” and “low” is ubiquitous in Western culture, it is by no means necessary. For instance, in Bali and Java pitches are conceived not as “high” and “low” but as “small” and “large.”²¹ Here the conceptual metaphor is PITCH RELATIONSHIPS ARE RELATIONSHIPS OF PHYSICAL SIZE, a mapping that reflects accurately the norms of acoustic production: small things typically vibrate more rapidly than large things. This acoustic fact is represented throughout the numerous parts of the gamelan, the collection of instruments central to the musical practice of Bali and Java.

Given that we structure our understanding of one domain in terms of another, what structures our understanding of the first domain? In 1987, Mark Johnson proposed that the process of cross-domain mapping was grounded in repeated patterns of bodily experience, which give rise to what he called image schemas.²² An image schema is a dynamic cognitive construct that functions somewhat like the abstract struc-

ture of an image, and thereby connects up a vast range of different experiences that manifest this same recurring structure.²³

As one example of an image schema, consider the VERTICALITY schema, which might be summarized by a diagram like that given in Figure 2. We grasp this structure repeatedly in thousands of perceptions and activities that we experience every day. Typical of these are the experiences of perceiving a tree, our felt sense of standing upright, the activity of climbing stairs, forming a mental image of a flagpole, and watching the level of water rise in the bathtub. The VERTICALITY schema is the abstract structure of the VERTICALITY experiences, images, and perceptions. Our concept of verticality is based on this schema, and this concept is in turn invoked by the various conceptual metaphors that use vertical space as a source domain through which to structure target domains such as emotions, consciousness, health, and musical pitch.

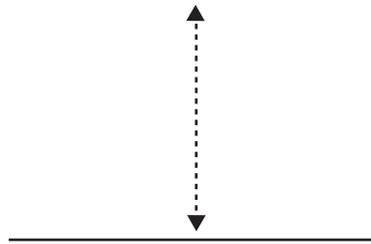


Figure 2: Diagram of VERTICALITY schema

The relationship between the VERTICALITY schema and our characterization of musical pitch in terms of an up-down spatial orientation is fairly immediate: when we make low sounds, our chest resonates; when we make high sounds, our chest no longer resonates in the same way, and the source of the sound seems located nearer our head. The “up” and “down” of musical pitch thus correlate with the spatial “up” and “down”—the vertical orientation—of our bodies. The VERTICALITY schema offers a straightforward way to explain why we characterize musical pitch in terms of high and low even when the actual spatial orientation of the means through which we produce pitches either does not reinforce the characterization or runs directly counter to it.

The theory of image schemata provides a way to explain how conceptual metaphors are grounded. However, it does not, by itself, explain why some conceptual metaphors seem intuitively better than others. To account for this, Lakoff and Turner proposed that such mappings are not about the imposition of the structure of the source domain on the target domain, but are instead about the establishment of correspondences between the two domains. These correspondences are not haphazard, but instead pre-

serve the image-schematic structure latent in each domain. Lakoff and Turner formalized this perspective with the Invariance Principle, which states that the components of the source and target domains involved in a cross-mapping should preserve the image-schematic structure of the target, and import as much image-schematic structure from the source as is consistent with that preservation.²⁴ Thus mapping the spatial orientation up-down onto pitch works because of correspondences between the image-schematic structure of components of the spatial and acoustical domains.

Let me now summarize a few of the points I have made about cross-domain mapping, and suggest that the view I have developed is still incomplete in one important respect. According to current theory, cross-domain mappings are grounded in repeated patterns of embodied experience called image schemata. These schemata provide the basic structure employed in the mappings: the VERTICALITY schema is thus fundamental to our understanding of two- or three-dimensional spaces as having directionality, and of musical pitch as “high” and “low.” Image schemata also constrain the possibilities for mapping between two domains, a constraint reflected in the Invariance Principle. Because the VERTICALITY schema can be applied to both the spatial and the musical domains, we can use our understanding of the former to structure our understanding of the latter. However, it remains to be explained why one mapping would be preferred over another—why, for instance, we tend to describe pitch relations in terms of “high” and “low” rather than “small” and “large”.

On the face of it, both of these mappings are equally viable. Both draw upon aspects of our embodied experience: on the one hand, countless experiences with the seeming origin of our own voices; on the other hand, countless experiences with the sounds given out by physical objects in the world around us. Both mappings allow us to describe musical pitches as elements within a continuum. And each mapping can be easily understood from the perspective provided by the other. For instance, in Camille Saint-Saëns *The Carnival of the Animals* the elephant’s part is taken by the contrabass (large is low), and the birds’ part by the flute (small is high).²⁵ Although musicians educated in the West sense the novelty of the mapping, they can nonetheless understand it perfectly well. In a similar fashion, musicians from Bali and Java, when confronted with Western conventions for notation, have few if any problems translating “small” and “large” into “high” and “low.”²⁶

The reason we prefer one mapping over another has to do with the global conceptual models we absorb from culture and that supply crucial support for the preferred mapping. In the West, the description of pitch relations in terms of “up” and “down” arose around the same time musicians began to develop ways of notating polyphonic compositions. These notational systems often relied, either directly or indirectly, on the physical placement of symbols on the page.²⁷ The attribution of “high” and “low” to musical

pitches is thus correlated with a system of notation that permitted both the visualization and preservation of musical works. This notation in turn relied on a global model that made three important assumptions: first, pitches could be regarded as objects that were independent of the sound source that produced them; second, graphical symbols could be used to represent these pitch-objects; and third, the surface on which these symbols appeared was analogous to physical space. In Bali and Java, the performance of music was associated with the rich palette of instruments through which it was effected. The attribution of “small” and “large” to pitches thus correlated with characteristics of the musical instruments intrinsic to musical performance. The conception of musical pitches as physical objects relies on a global model that does not, at some fundamental level, disassociate a pitch from the sound source that produces it.²⁸

As I have discussed elsewhere, the characterization of pitch relations can be informed by mappings other than “high” and “low,” and “small” and “large.”²⁹ In each case, the basic mapping relies on embodied knowledge and the correlation of the musical domain with a more concrete domain. The specific mapping chosen within a tradition of discourse about music reflects not so much absolute musical structure as it does the broader cultural practice within which music and its understanding are embedded: mappings reflect the conceptual models that are important to culture.

Cross-domain mapping, anchored by image schema theory and integrated into a broader account of the way culture informs our conceptualizations, offers a way to explain how musical concepts connect with those from other domains as part of a larger process of meaning construction. There is, nonetheless, one lingering problem: in the mappings I have discussed thus far, music appears to be the weak partner. Indeed, there is a notable asymmetry: we most typically map from some domain onto music, but not from music onto some other domain. However, there are circumstances in which musical concepts hold their own, and one of these involves combinations of music and text. Let us return to the opening of “Erlkönig” to explore this issue.

Conceptual Blending

One of the most immediately striking things about this opening is the way Schubert fashions his music to meet and to advance the narrative set out by Goethe’s text. One common explanation for how he accomplishes this relies on the notion of text painting: that is, Schubert’s music is effective because it imitates the sound of a horse galloping.³⁰ However, the extremely regular pattern of the right-hand figure is actually a poor imitation of the irregular sound made by a horse’s hooves at full gallop. Note also that almost a full fifteen measures pass before the voice enters and speaks of the act of riding—how can we understand the right-hand figure as even an idealized represen-

tation of a horse galloping before that moment? Simple mimesis, which is the basis for the conventional understanding of text painting, cannot account for the effectiveness of this music. Instead, what is at work is a coordinated set of mappings between Schubert's music and Goethe's text, mediated by our experiences with urgent missions and nightmarish scenes.

In order to explain similar mappings in language and literature, Mark Turner and Gilles Fauconnier developed the theory of conceptual blending. I shall not rehearse the theory of conceptual blending here; I will, however, offer one small extension to it. As originally presented by Turner and Fauconnier, a conceptual blend involves at least four mental spaces: a generic space, two input spaces, and the blended space. In most of the work on blending, the two input spaces are set up by language. In the following, I want to explore a conceptual blend in which one of the input spaces is set up by language, but the other space is set up by music. This is, of course, perfectly in keeping with mental space theory: as Fauconnier has noted, mental spaces are quite general cognitive structures that are set up for many purposes;³¹ language can set up mental spaces, but so can music. In the case of "Erlkönig," Schubert's music sets up one musical space, and Goethe's text another. As shown in the conceptual integration network (or CIN) diagrammed in Figure 3, these then serve as input spaces for a blend that combines concepts from both the musical and textual domains.

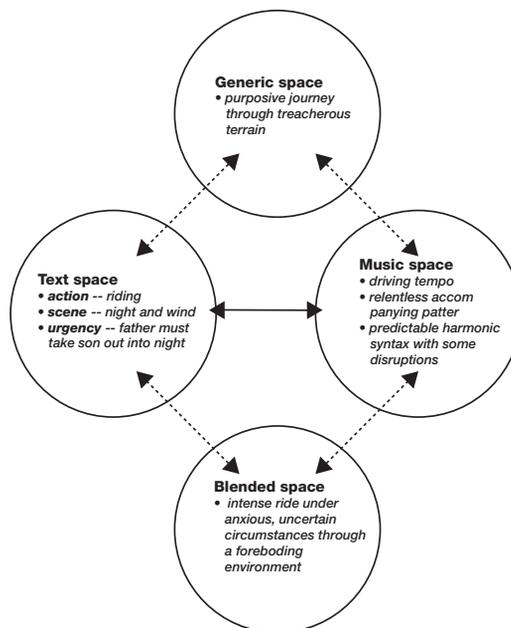


Figure 3: Conceptual integration network for text painting in Schubert's "Erlkönig"

As indicated, the text of the first stanza of the poem provides a basic description of action (namely, riding); sketchily sets the scene (night and wind); and suggests a somewhat diffuse sense of urgency and threat (circumstances require that the father take his son out into the night; during the ride the father holds the child safe and warm). The corresponding music establishes a driving tempo and relentless accompanimental pattern, begun *forte* and then (in measure-13) dropping to a *pianissimo*. Although the basic harmonic syntax is very clear, its predictability is offset somewhat by the bass motive. As I have noted, the motive sets up a two-bar hypermeter, which is confirmed by the contrast in articulation provided by the held notes of measures-6 through 7. The cadence of measures-7 through 8, however, complicates matters, for the two-octave Gs of measure-8 do dual duty as the completion of the cadence and as a repetition of measure-1. This elision tempts one to hear the two-octave Gs of measure-10 as similarly initiating a span of music—that is, as a beginning, rather than a conclusion. The situation becomes more complicated in measure 24, when the beginning note of the motive also marks the termination of the circle of fifths progression of measures-20 through 23. However, the low Bb of measure-25, instead of concluding the motive, marks the beginning of the next textual and musical phrase.

On the surface, the mental spaces set up by the text and the music seem to have few correspondences with each other. However, both text and music describe a directed progression through a relatively specific topography. In the textual space, the progression is through some sort of terrestrial nightscape, all the while impelled by an urgency of uncertain origin. In the musical space, the progression is through a portion of the tonal relationships of G minor, given life here as a sonic terrain marked by pounding rhythmic patterns, uncertainties in phrase structure, and sudden changes (such as appearance of the dominant of G minor in measure-29, or the return to *forte* in measure-32). There are thus a number of correlations between the two spaces, which can be understood to share a more abstract topography organized around the notion of a purposive journey through treacherous terrain—namely, the generic space shown in Figure 3. As usual with such blends, the emergent structure of the blended space of the this CIN can be described in terms of the operations of composition, completion, and elaboration.

Composition combines elements from the textual and musical spaces to yield new constructs in the blended space. The driving intensity of the music is linked to a ride through the night; the mercurial phrase structure of the music with the vaguely anxious circumstances of the ride and the buffeting of the night wind; the relatively dark timbre of the accompaniment (owing both to the predominance of the minor mode and thematic material in the bass register) with the nocturnal scene as a whole.

Completion fills out this picture. Although nowhere is a horse mentioned, the

urgency of text and music require nothing less than the strongest and swiftest animal available. Again, the rhythmic figuration is a poor mimicry of the irregular clatter of horses' hooves. However, the interlocking cycles of rhythmic activity (especially in the contrast between the repeated triplet-eighth figure and the phrase structure of the bass motive) are a perfect, if somewhat abstract, match for the interlocking cycles of physical motion between horse and rider at full gallop. Given the tempo of the music, the road over which father and son travel must be a clear one, rather than a twisting path through the forest. Finally, both music and text suggest that we join a ride already in progress.

Further elaboration of this image will, of course, be shaped by the continuation of the text and music. However, these introductory materials offer fertile ground for the imagination. Given the physical and psychic tumult sketched thus far, there must be more to the story: we simply must know what has occasioned and what threatens this midnight ride.

There is certainly more that could be said about text and music in Schubert's song. The perspective provided by conceptual blending has the potential to add to this discussion, for it gives a way to explain how music and text combine to create a rich world for the imagination. More importantly, it shows ways music, as an independent domain of discourse, can combine with discourse from other domains.

Musical Ontology

Before turning in earnest to musical syntax I want to take up one final problem specific to music, but that has implications that resound through our consideration of the relationship between cognitive science and the arts. The problem is that of musical ontology, and it is most clearly demonstrated by example. Shown in Example 2 is a score for Rodgers and Hammerstein's "My Favorite Things," one of the hits from their Broadway show *The Sound of Music*. This version of the tune is taken from one of the commercial fake books currently available. Such books offer minimal versions of upwards of five hundred popular songs—each consisting of the melody, words, and a straightforward harmonization rendered in shorthand chord symbols—allowing players familiar with the basic styles of popular music to "fake" their way through a performance if the tune is called during a professional engagement or a jam session.

Lively, with spirit

Em Cmaj7

Rain-drops on ros-es and whis-kers on kit-tens. Bright cop-per ket-tles and
Cream col-ored po-nies and crisp ap-ple stru-dels, Door-bells and sleigh-bells and

Am7 D9 G C

warm wool-en mit-tens, Brown pa-per pack-ag-es tied up with string,
schnitz-et with noo-dles, Wild geese that fly with the moon on their wings,

G C F#m7b5 B7 E

These are a few of My Fa-vor-ite Things. Girls in white

A

dress-es with blue sat-in sash-es, Snow-flakes that stay on my nose and eye-lash-es,

A6 D9 G C G C

SU-ver white win-ters that melt in-to springs, These are a few of My

F#m7b5 B7 Em F#m7b5 B7 Em

Fa-vor-ite Things. When the dog bites, When the bee stings, When I'm

C A7

feel-ing sad, I sim-ply re-mem-ber My Fa-vor-ite Things and

G C G C G D#9 D7 G D7 G

then I don't feel so bad.

Example 2: Score for “My Favorite Things” from a commercial fake book

Shown in Example 3 is a second score for “My Favorite Things,” taken from a jazz fake book. Jazz fake books are similar in most respects to commercial fake books, but they usually take as their point of reference performances by jazz musicians. While such books feature any number of popular tunes, they also include many pieces written by jazz musicians (as opposed to composers writing for Tin Pan Alley or the Broadway stage), all rendered in harmonies more typically used by jazz players. Thus the E-minor chord in Rodgers and Hammerstein’s original appears as an E-minor-seventh in the jazz fake book, and the C-major chord as a C-major-seventh. But there are other changes, both explicit and implicit: in the second measure of the jazz fake book version of the tune we find an F#-minor-seventh, something not found in the original. This change alludes to the locus classicus for rendering “My Favorite Things” in the jazz community, the performance recorded by John Coltrane in 1960. Joining Coltrane (who played soprano sax) on that session were McCoy Tyner (piano), Steve Davis (bass), and Elvin Jones (drums). The version they recorded is different still from the first two I have discussed. While the alternating E-minor-seventh/F#-minor-seventh found in the jazz

fake book is indeed prominent in the version by Coltrane’s group, they also adopt and significantly expand the two-measure interlude found in the original; they also cut the last two phrases of the tune—“When the dog bites, When the bee stings”—and simply cycle back to the alternating E-minor-seventh/F#-minor-seventh at the beginning of the tune. The recorded performance is a long one, over thirteen and a half minutes long, and resolutely hews to the format laid out in the opening minutes until the very end. At that point, the last two phrases of the original make their first and only appearance, and bring the tune to a close.

MY FAVORITE THINGS - RICHARD RODGERS

Example 3: Score for “My Favorite Things” from a jazz fake book

So here is the problem: which of these versions is “My Favorite Things”? If all them are equally good versions, what wouldn’t count as an example of “My Favorite Things”? If there are indeed “true” versions of the tune, how would we determine this truth?

Framed by these questions is the issue of musical ontology—that is, the ontological status of a work of music. In Western culture, it is common to regard pieces of music as having an objective status, speaking of a “work of music” as one might of a painting or sculpture. This habit of thought persists in spite of the transience of musical phenomena: unlike a painting or sculpture, a “musical work” has no enduring material existence. The enigma of “works” that are of profound cultural importance and yet are totally ephemeral has led to a number of attempts by philosophers and others to account for the objective status of the work of music and, by extension, music itself; thus the problem of musical ontology.³² A contrasting perspective has developed among ethnomusicologists who, working chiefly with non-Western music, have noted the fluidity of musical practice and a concomitant absence of any ascription of permanence to the products of musical activity. These observations render doubtful any universal solution to the problem of the ontology of the musical work, or even the ontology of “music,” since many cultures do not show linguistic or practical evidence of equivalent concepts.³³ And yet, musical practice within these cultures, while fluid, is certainly not unregulated. In general, the habit of objectifying aspects of musical practice is omnipresent; thus the problem of musical ontology, while most properly one that arises only among those who believe in an objective status for the work of music, also has applications to situations where no such object is posited.

My solution to this problem makes recourse to the theory of categorization I outlined above. In brief, knowing that a particular succession of sounds counts as an instance of “My Favorite Things” means knowing how to categorize those sounds in accordance with conceptual models shared with other members of a culture. More specifically, each “work of music”—or whatever other unit of cultural currency we choose to focus on—constitutes a natural category. On the one hand, the transparency and immediacy of the conceptual models that guide the process of categorization explain why we describe works of music in the same language we use for objects: we speak of “My Favorite Things” just as we speak of “the red cup,” even though the former (as sounding music) is a phenomenon completely unlike the latter.³⁴ On the other hand, because membership in natural categories is by degree, a certain amount of fluidity is possible, varying in degree with the cultural context relative to which the conceptual model for the category is framed. We might include a wide variety of performances in the category for “My Favorite Things”—some more, some less competent, perhaps even one done for a high school marching band—but could well draw the line at a version for pitched percussion and automobile horns.

As I have already suggested, shared conceptual models are the basis for the knowledge that is constitutive of culture. However, the sharing of conceptual models is by no means absolute: under certain circumstances only the basic outlines of a model may be shared, with individuals realizing the details in any number of disparate ways. Should the differences in realizations become profound, disagreements about behavior or breakdowns in communication may occur. And if such problems persist, features of the model may be opened up for negotiation and change. Models also change because of the inherent creativity of humans, and because the conceptual models proper to one domain of culture may be shared with or influenced by models from other domains. In sum, we should not only expect that a variety of sound-sequences will be accepted as renderings of “My Favorite Things,” but that the range of what is acceptable will vary according to the cultural framework relative to which categorical judgments are made.

To illustrate this approach to how determinations of musical identity are made, let me propose a conceptual model for “My Favorite Things” focused on its earliest version, that is, as a tune from a Broadway show or, more simply, a popular tune. The model, which is diagrammed in Figure 4, consists of five correlated conceptual elements, characterized as the things necessary for an adequate performance of the song. Given this model, the first version of “My Favorite Things” would be close to the center of the category. The second version, slower, lacking not only the words but the two-measure interlude, and making use of improvisation—is rather less typical of the category. Finally, Coltrane’s version—with its loping groove, omission of the final two phrases of the song, and fragmentation of the remaining phrases to refract, rather than simply illuminate, their constituent musical materials—Coltrane’s version is close enough to the boundary of the category that someone obsessed with the ‘original’ version of the tune might refuse to acknowledge it as an acceptable version of “My Favorite Things.” Of course, were we to take Coltrane’s rendering of “My Favorite Things” as our model for the category, our evaluation of the different versions would be exactly reversed: it would now be Mary Martin or Julie Andrews who was out in the cold.

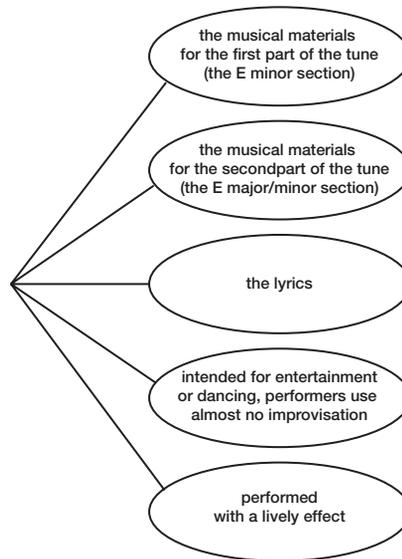


Figure 4: Conceptual model for Broadway show version of "My Favorite Things"

Perhaps most intriguing is the way Coltrane's version changed irrevocably how musicians regard "My Favorite Things." Despite the fact that Benny Goodman and his orchestra recorded most of the songs from *The Sound of Music* in the same year Coltrane's recording came out, it was only "My Favorite Things" that entered the jazz repertoire, and it was in Coltrane's version, not Goodman's. This change points out something important about tunes within the popular performance tradition: they are subject to change over time. This change can be a consequence of commercial pressures—the need to create new commodities to replace those unsuccessful, worn out by too frequent use, or merely branded as "old"—but it can also be a consequence of innovation of the sort behind Coltrane's re-invention of Rodgers and Hammerstein's tune. Through his performance, Coltrane offered a new conceptual model for "My Favorite Things."

Given these examples, it would seem that musical identity (determinations of which are made relative to conceptual models) is necessarily fluid rather than fixed. While this may be true over the long run, in practice what counts as an instance of a musical work changes hardly at all. This is because the conceptual models relative to which determinations of musical identity are made are part of cultural knowledge: they are typically shared among the members of a musical community. This is not to say that the stability of this cultural knowledge cannot be challenged. In the case of musical practice, such challenges might come from a number of sources: from the innovations of influential individuals, from contact with other cultures, from a sudden increase in the complexity

of musical practice, or from changes in technology. Faced with this sort of challenge, a musical community may develop strategies to eliminate or control the instability of the conceptual models upon which the practice of that community is based.

One common strategy is to represent key structural elements of a local conceptual model through artifacts such as scores or recordings. Although the score illustrations used for “My Favorite Things” carry the aura of authority, they actually follow from the various conceptual models for the song: a score is an artifactual manifestation of the elements of the conceptual model deemed most relevant to the musical practice of which the model is a part, created as a means of stabilizing the model.

Viewed from the perspective of categorization, our sense that there is a work of music and our acceptance of multiple and oftentimes incommensurate versions of this work are both unremarkable. If we think of “the work of music” as a natural category, our tendency to ascribe objecthood to it can be traced to the way conceptual models simplify our perspective on ontology by representing the whole of a category through one idealized type. The category “My Favorite Things” is like the category cat: simple and unitary, complex only when we start to tease out all of its ramifications. However, the conceptual model not only gives us a simplified cognitive representation of the diverse members of a category, but it provides a tool for reasoning about them. Thus we can explain why a tiger is a less good member of the category cat than is a tabby, and why a score, recording, or performance of “My Favorite Things” that leaves out the melody, alters the chords, and changes the form is a less good member of the category “My Favorite Things” than one that does not.

My consideration of musical ontology has, of necessity, focused on music. Nonetheless, I believe the approach I have outlined has applicability to other expressive media. The applications to performed media, such as dance and the theater, are relatively straightforward. The application to artifactual media, such as painting or sculpture, would seem rather more strained: what would it mean to approach Michelangelo’s David as a cognitive category organized around a conceptual model? And yet we know our regard of expressive media, both performed and artifactual, changes over time, suggesting that music, rather than being the exceptional, is instead thoroughly typical of the things that make up our aesthetic universe.

Toward a Cognitive Grammar of Music: The Blues

I would now like to take the perspective I have developed thus far and move in an even more adventurous direction: I would like to show how it can contribute to the development of a cognitive grammar of music. As will become evident, I adopt the view of

some cognitive linguists and assume that syntax and semantics are poles on either end of a continuum—the notion of grammar, as I use it, absorbs both. Further, I assume (along with some cognitive linguists) that syntactic structure has its basis in cognitive processes like categorization and cross-domain mapping. Although I have used this approach to analyze syntactic structure in music by composers like Wolfgang Amadeus Mozart, Richard Wagner, and Johannes Brahms, in what follows I shall explore aspects of musical syntax through an account of the grammar of the blues.

As a relatively unproblematic example of the blues, I would like to focus on a blues recorded in 1974 by Henry Roeland Byrd, known throughout his performing career as Professor Longhair; the title of the song is “They Call Me Doctor Professor Longhair.” This particular blues is what is conventionally called a twelve-bar blues, so called because it is based on a repeating cycle of twelve measures, or bars. Since an understanding of this pattern will be important for what follows, I would like to explore its structure in just a bit of detail.

Example 4 is a simplified representation of one version of the harmonic structure of a twelve-bar blues. In this version, measures 1 through 4 are occupied exclusively with tonic harmony (F7). In measures 5 and 6 there is a temporary shift to the subdominant (Bb7), but tonic returns once more in measures 7 and 8. The rate of harmonic change then undergoes a decided acceleration, for there follows at one-measure intervals dominant (C7), subdominant, tonic, and finally dominant once more.

The image shows two staves of musical notation in 4/4 time, representing the harmonic structure of a twelve-bar blues. The first staff contains measures 1 through 8. Measures 1-4 are F7. Measures 5-6 are Bb7. Measures 7-8 are F7. The second staff contains measures 9 through 13. Measure 9 is C7. Measure 10 is C7. Measure 11 is Bb7. Measure 12 is F7. Measure 13 is C7. The notation uses chord symbols above the staves and chord diagrams below the staves.

Example 4: Basic harmonic pattern for a twelve-bar blues

The basic rhythmic structure of the blues is a resolute four beats per measure, but there is considerable nuance in the rhythmic pattern itself. In the first case, each of these beats is typically divided into three parts, as shown in Example 5, yielding what is often called a “12/8 feel.” In some cases these divisions are stated explicitly, sometimes they simmer just below the surface of the music, but in either case they impart a rolling feel to the basic beat. While the first and third of the main beats of each measure each get an accent (as is typical for quadruple meters), the second and fourth take a competing accent, which is usually called a backbeat.



Example 5: Basic rhythmic pattern for a twelve-bar blues

A third important pattern within the blues is that of the lyrics. As can be seen, these are in an AAB pattern, with each line of text extending over four measures of music.

- A Well, all the boys call me Dr. Professor Longhair,
the girls all call me little ol' lovin' man,
- A All the boys call me Dr. Professor Longhair,
but the girls all call me little ol' lovin' man,
- B Well they know I'm not a doctor, baby,
they know I'll give it (indistinct) [all I can]

Although the second line, as a repetition of the first, would seem to do little to advance the narrative of the lyrics, it actually has three important semantic functions. First, it serves to sustain the mental space set up by the first line, which gives the listener a chance to survey that space in just a bit more detail. Second, it gives the singer an opportunity to re-organize the space by embellishing the words, or by changing their delivery to draw out alternate meanings. Third, the repeated line is coordinated with the first significant change in harmony, which has the effect of further emphasizing the polysemous nature of the words.

These three elements are basic to the syntactic structure of a twelve-bar blues. But before proceeding to a fuller account of that syntax, I want to consider briefly the relationship between syntax and time. Music is often regarded as the temporal medium par excellence, but when we start to think about syntactic structure the emphasis on music's temporality can seem a bit odd. If syntax is concerned with the way materials are ordered so as to make communication possible, then the syntax of every expressive medium has some temporal dimension, if only as "first" and "second," or "before" and "after." What makes music special is that it typically involves multiple levels of syntax, and much of musical meaning relies on how these levels are coordinated.

In the blues this coordination begins with two syntactic levels—one rhythmic, the

other harmonic—which establishes a framework for meaning construction. Rhythmic syntax is based on incidence and duration relative to some temporal frame. At its most primitive—for instance, in recitative, or in unaccompanied chant—this temporal frame is practically the same as that for language. Where this frame is made up of regularly recurring accentual cycles, as it is in the blues, it is metric. Harmonic syntax is based on pitch relationships that are established by pitch assemblies; the chords of Example 4 are examples of such pitch assemblies. Syntactic relationships at the harmonic level are to other pitches or pitch assemblies. Such relationships can be relatively concrete: for instance, note that in Example 4 the tonic and the subdominant share one pitch (F), and the tonic and dominant share one pitch (C), but the subdominant and dominant share no pitches. Pitch relationships can also be comparatively abstract: for example, all of the chords of Example 4 are of the same type, meaning that they are transpositionally related to one another.

Let us now consider specific aspects of each of these syntactic levels as manifested in a typical twelve-bar blues. The basic syntax of the rhythmic level is fairly straightforward: a 12/8 feel is set up at the beginning, and then continues for the duration of the piece. Added to this are important articulations: each four measures combines to create a phrase (something that is often marked by an extra push at the beginning of measures 1, 5, and 9). Further, measure 12, which brings the form to a close, is also typically given a clear articulation to mark the point when the form turns back around to the beginning. This syntax can be characterized in cognitive terms with reference to two structures: first, categories of rhythmic patterns can be discerned at the level of the measure, the phrase, and the form as a whole; second, the rhythmic pattern correlates with the bodily motion of dancing—as Simon Frith has noted, the rhythm of popular music (including the blues) is the rhythm of dance.³⁵

The harmonic level of syntax is, for the most part, occupied with strategies that privilege the tonic. Tonic is the only harmony heard in the first four measures. The next four measures begin with the contrasting pitch assembly of the subdominant, but this is then followed by a return to the tonic. The final four measures start with a further departure, this time to the dominant; return briefly to the subdominant; and then once again arrive on the tonic, which is followed by a final contrasting dominant. This syntax can be characterized in cognitive terms with reference to a cross-domain mapping from our experience with various journeys, or even bodily movements. These begin in one place (tonic), go to another (subdominant) return to the starting point (tonic), go to yet another (dominant) and then, after passing back through the first site visited (subdominant) return ultimately to the starting place (tonic), the whole process summarized by a final reference to the departure that signaled the potential for these various movements to continue *ad infinitum* (that is, the dominant).

These two syntactic levels connect at a number of points. First and foremost, as demonstrated by Professor Longhair's performance, the rhythmic level can be projected almost exclusively through the pitches that make up the harmonic level. In general, the harmonic frame also reinforces the phrase structure of the rhythmic frame (as does the structure of the lyrics). And the rhythmic frame provides support for embellishments to the harmonic frame, as when non-chord tones are added on the backbeat. Despite these various overlaps, I would like to continue to regard these levels as functionally independent: we do not have to have the harmonic level present for the rhythmic level to be present, and vice versa.

In the blues, these two coordinated levels provide a framework for the melody, the one syntactic feature of the genre that I have not discussed thus far. The harmonic syntax for blues melodies is organized around a five-note pitch assembly called a minor pentatonic scale, which is shown in Example 6a. This scale can also be expanded to seven notes, yielding the transposed Dorian mode shown in Example 6b. Melodies created from these pitch assemblies are then sung or played over the basic syntactic framework provided by the rhythmic and harmonic levels.



Example 6: Basic melodic materials for a twelve-bar blues

One thing striking about this arrangement is that the pitch assemblies of Example 4 and Example 6 do not, in all cases, match. Note in particular that the thirds of the tonic (F7) and dominant chords (C7)—A-natural and E-natural respectively—conflict with the Ab and Eb of Example 6. This conflict is the source of what are called “blue notes,” and points to the strategies through which meaning is constructed in blues performance.

The basic strategies for meaning construction in the blues are two. First, the performer can conform with the background structure; remember that this structure is itself complex, consisting of both rhythmic and harmonic syntactic levels. On the other extreme, the performer could depart totally from the background structure. In practice, performers move back and forth between these two strategies to shape their musical discourse. The simplest analogue, and one in terms of which listeners often hear such discourse, is to relative distance within a visual field, where syntactic conformance is associated with objects that are close to the viewer, and syntactic non-conformance is associated with objects that are distant from the viewer.

Now there is much, much more that could be said about meaning construction in the blues. I have deliberately neglected the words to the blues, as well as the cultural and social context within which blues performance takes place. But I hope the glimpse I have provided of how one might describe, in cognitive and practical terms, the grammar of the blues gives some sense of the potential of applying research in cognitive science not only to music, but to art as a whole.

Summary

Music has been important to the anthropologist Claude Lévi-Strauss throughout his long career, not the least because its mysteries symbolized to him all that was powerful and obscure in human experience. Perhaps because of its importance to him he placed it in an exclusive world sealed off from most of humanity. In introductory comments in *The Raw and the Cooked* he contrasts the mystery of music with that of mythology:

Music raises a much more difficult problem, because we know nothing of the mental conditions in which musical creation takes place. In other words, we do not understand the difference between the very few minds that secrete music and the vast numbers in which the phenomenon does not take place, although they are usually sensitive to music. However, the difference is so obvious, and is noticeable at so early an age, that we cannot but suspect that it implies the existence of very special and deep-seated properties. But since music is a language with some meaning at least for the immense majority of mankind, although only a tiny minority of people are capable of formulating a meaning in it, and since it is the only language with the contradictory attributes of being at once intelligible and untranslatable, the musical creator is a being comparable to the gods, and music itself the supreme mystery of the science of man, a mystery that all the various disciplines come up against and which holds the key to their progress.³⁶

As is evident from what I have presented to you today, I do not share Lévi-Strauss's view that musical expression is an exclusive domain whose inner secrets are accessible only to a chosen few. Both the creation and the appreciation of music involves cognitive processes common to all humans, yielding the categories, cross-domain mappings, and conceptual models through which we structure our understanding of the world. This is not to say that I think that all humans have an equal capacity for making or appreciating music—for there are the truly gifted as well as the truly obtuse among us—but only that the resources we draw upon to create musical understanding are all fundamentally the same. The exclusivity noted by Lévi-Strauss is as much a product of an entrenched nineteenth-century idea of musical value as it is an assessment of the profound mysteries of music. I do, however, agree with Lévi-Strauss that music does have a lot to do with the mystery of what it means to be human, and that it offers clues to our understanding of what that means.

Music is a thoroughly human activity—indeed, given what we now know about

human cognitive structure, I'm not quite sure why it would be regarded as anything else—and if we are to understand it we must explore the cognitive processes through which humans structure their understanding of the world as a whole. I do not for a moment imagine that this will be easy, or that finding answers to the questions we ask will not take some time and trouble. But I believe the tools are there, if we only set our hands to them. By this means we can fully acknowledge what is evident in our many encounters with music: that works like Schubert's "Erlkönig," Rodgers and Hammerstein's "My Favorite Things," Professor Longhair's eponymous blues, and many, many others besides, have meaning for us because they connect most deeply with what it means to be human.

NOTES

1. This paper was originally presented at the Winter Symposium "Art and Cognition—Analyses and Perspectives," held January 25-27, 2001 at the Center for Semiotics, University of Aarhus. The present version is somewhat shorter, and (of necessity) omits the musical examples that were played during the original presentation.
2. Mark Turner, "Design for a Theory of Meaning," in *The Nature and Ontogenesis of Meaning*, edited by Willis F. Overton and David S. Palermo (Hillsdale, New Jersey: Lawrence Erlbaum Associates, Publishers, 1994), 91-92.
3. Antonio R. Damasio, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness* (New York: Harcourt Brace & Company, 1999), 320.
4. My approach to ways recent work in cognitive science can be applied to problems in musical understanding can be found in *Conceptualizing Music: Cognitive Structure, Theory, and Analysis* (New York: Oxford University Press, in press). This essay draws upon and expands the research presented there.
5. Ray Jackendoff, *Consciousness and the Computational Mind* (Cambridge, Massachusetts: MIT Press, 1987); Mark DeBellis, *Music and Conceptualization* (Cambridge: Cambridge University Press, 1995).
6. I have argued this in more length in my essay "Seeger's Unitary Field Theory Reconsidered," in *Foundations of Modern Musicology: Understanding Charles Seeger*, edited by Bell Yung and Helen Rees (University of Illinois Press, 1999), 130-49; the point is pursued further in the next section.
7. See, for instance, Gerald M. Edelman, *The Remembered Present: A Biological Theory of Consciousness* (New York: Basic Books, 1989); James Hampton and Danièle Dubois, "Psychological Models of Concepts," in *Categories and Concepts: Theoretical Views and Inductive Analysis*, edited by Iven van Mechelen, James Hampton, Ryszard S. Michalski, and Peter Theuns, *Cognitive Science Series* (London: Academic Press, 1993), 11-33; Lawrence W. Barsalou, "Flexibility, Structure, and Linguistic Vagary in Concepts: Manifestations of a Compositional System of Perceptual Symbols," in *Theories of Memory*, edited by A. F. Collins, S. E. Gathercole, M. A. Conway, and P. E. Morris (Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1993), 29-101; Lawrence W. Barsalou, et al., "Concepts and Meaning," in *Chicago Linguistics Society 29: Papers from the Parasession on Conceptual Representations*, edited by K. Beals, et al. (Chicago: University of Chicago, Chicago Linguistics Society, 1993); Edward E. Smith and Douglas L. Medin, *Categories and Concepts* (Cambridge, Massachusetts: Harvard University Press, 1981); Gregory L. Murphy and Douglas L. Medin, "The Role of Theories in Conceptual Coherence,"

- Psychological Review* 92, no. 3 (July 1985): 289-316. I have developed the notion of a musical concept in my essay "Seeger's Unitary Field Theory Reconsidered."
8. Perhaps the best support for the independence of categorization from language comes from animal research. See, for instance, Donald R. Griffin, *Animal Minds* (Chicago: University of Chicago Press, 1992).
 9. Lawrence W. Barsalou, "The Instability of Graded Structure: Implications for the Nature of Concepts," in *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, edited by Ulric Neisser (Cambridge: Cambridge University Press, 1987), 101-40.
 10. Eleanor Rosch, "On the Internal Structure of Perceptual and Semantic Categories," in *Cognitive Development and the Acquisition of Language*, edited by Timothy E. Moore (New York: Academic Press, 1973), 111-44; Eleanor Rosch, "Cognitive Representations of Semantic Categories," *Journal of Experimental Psychology: General* 104, no. 3 (1975): 192-233.
 11. See also the discussion in my *Conceptualizing Music: Cognitive Structure, Theory, and Analysis* (New York: Oxford University Press, 2001), chapter 2.
 12. My thinking about conceptual models has been shaped in important ways by research in artificial intelligence from the 1970s and 1980s—in particular, the work of Marvin Minsky, and Roger Schank and Robert Abelson—and I find it useful to characterize conceptual models in terms of simple propositions linked together to create a very basic system of inference. My representations of conceptual models as a collection of interlinked nodes reflects this influence. For more on categories and conceptual models see *Conceptualizing Music*, chaps. 2 and 3.
 13. Describing Schubert's compositional strategies in terms of a process of departure and return relies on a cross-domain mapping that accounts for relationships between musical events in terms of a journey. The application of cross-domain mapping to music is taken up in the next main section below.
 14. For a recent and stimulating sally in this debate, see Jerry A. Fodor, *Concepts: Where Cognitive Science Went Wrong* (Oxford: Clarendon Press, 1998).
 15. Edelman, *The Remembered Present*, 141.
 16. Others who either make a strong connection between categories and concepts or assert their equivalence include Hampton and Dubois, "Psychological Models of Concepts.," Barsalou, "Flexibility, Structure, and Linguistic Vagary in Concepts.," Barsalou, et al., "Concepts and Meaning.," Smith and Medin, Categories and Concepts; Murphy and Medin, "The Role of Theories in Conceptual Coherence."
 17. Griffin, *Animal Minds*, chapter 6. A similar approach can be seen in the work of Douglas Hofstadter and his associates, which has resulted in a number of computer programs that deal with relatively compact bundles of information they call "concepts;" see Douglas R. Hofstadter and the Fluid Analogies Research Group, *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought* (New York: Basic Books, 1995).
 18. The idea that language was not required for musical analysis was one of the ideas behind Hans Keller's wordless analyses. See Hans Keller, "Wordless Functional Analysis no. 1: Mozart, K. 421," *The Score* 22 (1958): 56-64; idem, "The Musical Analysis of Music," in *Essays on Music*, ed. Christopher Winzle (Cambridge: Cambridge University Press, 1994), 126-28; Keller, "Functional Analysis no. 9a: Mozart's Piano Sonata in A Minor, K. 310," in *Essays on Music*, 129-38.
 19. Jackendoff, *Consciousness and the Computational Mind*; DeBellis, *Music and Conceptualization*.
 20. I have discussed this at length in my essay "Metaphor and Music Theory: Reflections from Cognitive Science," *Music Theory Online* 4, no. 1 (January 1998); see also *Conceptualizing Music*, chapter 2.

21. Benjamin Brinner, personal communication, 8 July 1997. See also Wim van Zanten, "The Tone Material of the Kacapi in Tembang Sunda in West Java," *Ethnomusicology* 30 (1986): 85.
22. Mark Johnson, *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. (Chicago: University of Chicago Press, 1987).
23. Mark Johnson, *The Body in the Mind*, 2.
24. Mark Turner, "Aspects of the Invariance Hypothesis," *Cognitive Linguistics* 1, no. 2 (1990): 254; emphasized in original. For additional writings on the Invariance Principle (which at first was called the Invariance Hypothesis) see George Lakoff, "The Invariance Hypothesis: Is Abstract Reason Based on Image-Schemas?" *Cognitive Linguistics* 1, no. 1 (1990): 39-74; Mark Turner, "An Image-Schematic Constraint on Metaphor," in *Conceptualizations and Mental Processing in Language*, edited by Richard A. Geiger and Brygida Rudzka-Ostyn, *Cognitive Linguistics Research*, vol. 3 (Berlin: Mouton de Gruyter, 1993), 291-306; and Mark Turner, *The Literary Mind* (New York: Oxford University Press, 1996), chapter 3.
A preliminary discussion of a similar sort of topographical invariance, with applications to music, can be found in Peter Gärdenfors, "Semantics, Conceptual Spaces and the Dimensions of Music," in *Essays on the Philosophy of Music*, edited by Veikko Rantala, Lewis Rowell, and Eero Tarasti, *Acta Philosophica Fennica*, vol. 43 (Helsinki: Philosophical Society of Finland, 1988), 9-27.
25. Camille Saint-Saëns, *Le Carnaval des animaux: Grande fantaisie zoologique*, ed. Felix Aprahamian (Zurich: Eulenberg, 1974), "L'éléphant" (no. 5), p. 11, "Volière" (no. 10), pp. 23-27.
26. Personal communication from Sumarsam, 10 April 1998.
27. With regard to an indirect reliance on physical placement, it was often a part of the pedagogy of daseian notation (developed in the ninth century, apparently with an eye toward accommodating the needs of organum at the fourth) to place "high" notes higher on the page, even though no such discrimination was required by the notation (since it relied strictly on a limited repertoire of letter-like symbols to indicate different pitches). See Hans Schmid, ed., *Musica et Scolica enchiridis una cum aliquibus tratatulis adiunctis*, *Veröffentlichungen der Musikhistorischen Kommission*, Bayerische Akademie der Wissenschaften, vol. 3 (Munich: Verlag der Bayerische Akademie der Wissenschaften, 1981); and the *Musica enchiridis and Scolica enchiridis*, edited by Claude V. Palisca, translated by Raymond Erickson, Music Theory Translation Series (New Haven: Yale University Press, 1995). These treatises are currently thought to have been written no later than the tenth century and possibly as early as the middle of the ninth. For further discussion of the history of this notational convention, see Marie-Elisabeth Duchez, "La représentation spatio-verticale du caractère musical grave-aigu et l'élaboration de la notion de hauteur de son dans la conscience musicale occidentale," *Acta musicologica* 51, no. 1 (1979): 54-73.
28. It should be observed that a disassociation of pitch from sound does not obviate an image-schematic basis for understanding pitch relations. It only bears witness to a reification of pitch subsequent to its initial conceptualization.
29. See *Conceptualizing Music*, chap. 2.
30. This understanding of text painting draws strongly on mimetic theories of musical meaning, which were particularly prominent in the eighteenth century. For discussions of such theories, see John Neubauer, *The Emancipation of Music from Language: Departure from Mimesis in Eighteenth-Century Aesthetics* (New Haven: Yale University Press, 1986), Downing A. Thomas, *Music and the Origins of Language: Theories from the French Enlightenment*, *New Perspectives in Music History and Criticism* (Cambridge: Cambridge University Press, 1995); and James H. Johnson, *Listening in Paris: A Cultural History* (Berkeley: University of California Press, 1995), chapter 2.

31. Personal communication from Gilles Fauconnier. This perspective is also represented, somewhat less explicitly, in Gilles Fauconnier and Mark Turner, "Conceptual Integration Networks," *Cognitive Science* 22, no. 2 (April-June 1998): 133-87 and Gilles Fauconnier, *Mappings in Thought and Language* (Cambridge: Cambridge University Press, 1997).
32. Thoughtful discussions of the problem of musical ontology can be found in Lydia Goehr, *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music* (Oxford: Clarendon Press, 1992); Philip V. Bohlman, "Ontologies of Music," in *Rethinking Music*, edited by Nicholas Cook and Mark Everist (Oxford: Oxford University Press, 1999), 17-34; and Nicholas Cook, "At the Borders of Musical Identity: Schenker, Corelli and the Graces," *Music Analysis* 18, no. 2 (July 1999): 179-233.
33. Examples of cultures the languages of which lack a word equivalent to "music" or concepts equivalent to that of a "musical work" include the Hausa of Nigeria, the Macuma Shuar of Ecuador, and the Mapuche of Argentina; Charles Keil reports such a lack in over twenty languages of the African continent and gives detailed discussion of the methodological problems involved in studying such cultures. See David W. Ames and Anthony V. King, *Glossary of Hausa Music and Its Social Contexts* (Evanston, Illinois: Northwestern University Press, 1971), ix; William Belzner, "Music, Modernization, and Westernization Among the Macuma Shuar," in *Cultural Transformations and Ethnicity in Modern Ecuador*, edited by Norman E. Whitten, Jr. (Urbana, Illinois: University of Illinois Press, 1981), 735-36; Carol E. Robertson-DeCarbo, "Tayil as Category and Communication Among the Argentine Mapuche: A Methodological Suggestion," *Yearbook of the International Folk Music Council* 8 (1976): 39; and Charles Keil, *Tiv Song* (Chicago: University of Chicago Press, 1979), 27-52.
I should make clear that what I find dubious is a universal solution to the problem of the ontology of a musical work. I do not think it is impossible, however, to speak about "music" in the case of a culture that does not show linguistic or practical evidence of such a concept, as long as one does not attempt to colonize that culture by insisting that its members, when they try to make sense of their own experiences and cultural practices, make use of the term.
34. There are, of course, aesthetic perspectives that support this habit (such as that associated with the "work-concept," as described by Lydia Goehr in *The Imaginary Museum of Musical Works*, 89-119) and technological innovations that give it credence (like music notation and sound recording), but these are not preconditions for the notion of a musical object.
35. Simon Frith, *Performing Rites: On the Value of Popular Music* (Cambridge, Massachusetts: Harvard University Press, 1996), chapter 6.
36. Claude Lévi-Strauss, *The Raw and the Cooked: Introduction to a Science of Mythology*, Volume 1, John Weightman and Doreen Weightman (Chicago: University of Chicago Press, 1969), 18.