Music, syntax, and the meaning of "meaning"

Geraint A. Wiggins
School of Artificial Intelligence
Division of Informatics, University of Edinburgh
80 South Bridge, Edinburgh GB-EH1 1HN, Scotland

Keywords: Computational Linguistics; Musical Cognition, Analysis and Composition

Abstract

I discuss the issue of meaning, and the definition of "meaning" in music. I propose that it is a mistake to import the linguistic notion of semantics into a musical context on the grounds that musical communication serves a different function and is of a different nature from linguistic communication, and that there is no *evidence* to support the suggestion that the two should function in a strongly similar way.

1 Introduction

"What is the meaning in music?", and the preceding issue "*Is* there meaning in music?", are questions which have been discussed in the world of musicology and, latterly, the cognitive sciences of music, at some length. One apparently natural means of approach to the problem is by means of devices borrowed from general linguistics, such as grammatical analysis and the construction of corresponding (compositional) systems of meaning.

In this paper, I argue that, while there are indeed many useful ideas to be borrowed from linguistics and logicians in the notions of syntax and semantics and their respective analyses, there are two inherent dangers of a naïve, wholesale importation: first, of confusion, on account of lack of clear definition of terms; and second, of a tendency to make the foot fit the shoe, rather than *vice versa*.

First, I present a summary of a few pieces of work in the area of grammatical approaches to the representation of musical structure, and discuss their contributions to the debate on the nature of musical "meaning". Then I introduce three different current definitions of "semantics", and go on to argue against the over-literal interpretation of linguistic-style semantics in a musical context; I propose an alternative terminology to avoid and highlight awareness of the problem.

Finally, by means of a simple example, I explain why the syntax/semantics dichotomy is important to (computational) linguists, and then argue that it simply does not apply to the musical context.

2 Musical Structure and Grammar

2.1 Why should grammars be useful for music?

Baroni et al. (1984) give a useful list of five features of natural languages which suggest that the transfer of methods of linguistic description to that of music will be fruitful. They may be summarised as follows:

Semiotics Music is primarily concerned with communication. So organised sounds are *signs of something else.*

Conventionality Accepted social conventions *ascribe a sense* to the form.

Variation Accepted structures vary between cultures and with time.

Structure Music is structured, and different elements may be associated with different functions.

Hierarchy Music is hierarchical in form.

This reasoning makes explicit many of the implicit assumptions in much of the work done on musical grammars, and serves as a useful starting point for this discussion.

Of particular interest in the context of the current argument is the point about *conventionality*, where "sense" is ascribed. "Sense" is yet another word which is often (at least in a linguistic context) used more or less interchangeably with "meaning". Meyer (1956), along with many others, assumes that "musical meaning" is a term which can be used without fear of confusion, which I suggest is not the case. I will return to this point in section 3; for the moment, the use of *grammars* is the focus of this discussion.

2.2 Musical grammars

There have been numerous attempts to describe music in more or less grammatical terms, such as those of Lerdahl and Jackendoff (1983), Steedman (1996), Bel and Kippen (1992) and Ponsford et al. (FC). These four examples alone show four very clearly disparate reasons for and/or approaches to using grammatical formalisms.

Lerdahl and Jackendoff attempt to explain the results of the cognitive processes involved in the mental analysis of (Western tonal) music, by means of hierarchical structures, specified in terms of four classes of rules which allow notes and groups of notes to be grouped together in what are intended to be cognitively significant ways.

Steedman's interest is related, but more specific: capturing the mental processes which lead to expectation (and hence acceptability) in Jazz progressions. Steedman's rules are much more like what a linguist would call a grammar, though he is forced by the restricted nature of the structures he is trying to capture to introduce "conventions": implicit meta-rules governing the relationships between different parts of the structures his grammar generates.

Bel and Kippen's work is less explicitly cognitively motivated: they aim to reproduce a style of Indian tabla drumming – the *qa'ida* – and to do so, not unsuccessfully according to their empirical studies, they use a so-called *pattern grammar* (Angluin 1980). Pattern grammars are a specialised grammatical formalism which extend standard Chomskian grammars (Chomsky 1965), to cover the kinds of meta-level restrictions stated as "conventions" by Steedman.

Finally, the work of Ponsford et al. is entirely an experiment to see how far a simple formalism intended for learning linguistic grammars can be taken when applied to capturing musical style.

The idea common to all of these approaches, and to the other work in similar vein going back as far as Schenker (1992, for example), is that music can usefully be described hierarchically[1], and, in particular, a grammatical formalism may be used to give a finite (and therefore manageable) description of an infinite (and therefore intractable) set of structures. It is considered to be a good idea to apply such approaches because the essential intuition that *structure* makes music what it is is so fundamental as to be universally agreed.

Some of the work mentioned above, however, has already raised an objection to the naïve import of standard grammatical technology to deal with musical structure. As Bel and Kippen point out, music bears a much clearer relation to poetry than it does to arbitrary spoken or written text, and as such might indeed be expected to fall under similar structural constraints.

This issue aside, syntax is still not usually enough. An English grammar, with no accompanying notion of meaning, will produce engaging but pointless statements such as Chomsky's famous

> "Colourless green ideas sleep furiously".

In Steedman's terms, this might be like generating an archetypal 12-bar blues – with absolutely no variation, and hence no interest at all. However, Bel and Kippen seem to get by perfectly well with no distinct representation of, or reference to, meaning, and are able to generate empirically acceptable *qa'ida* pieces with their pattern grammar alone; equally, Ponsford et al. produce a reasonable approximation of a given musical style without recourse to meaning.

2.3 What is a grammar for?

It is important to add here that the music research world is not unanimous in its acceptance of the quasi-linguistic, syntactic approach. For example, Dempster (1998) neatly sums up his argument against the grammatical assumption as follows:

> "... while music typically has very elaborate and regular structures – much like language – these structures do not apparently originate from nor are they in the service of the need to encode meanings – exactly unlike language."

This contains, of course, a strong claim about musical meaning, to which I will return in section 3, below.

The point for the purposes of the current section is this: the attempt at refutation quoted above appears to be founded on an invalid inference. The argument presented is that (even though it has significant structure) music does not have meaning, and *therefore* there is not a grammar for it. Postponing until later the issue of meaning, the invalid inference is that because something does not have meaning it is not possible (or perhaps appropriate?) to describe its structure with a grammar. This is demonstrably false, the work of Mitchell (1992), on architecture, being a good counter-example[2]. The work of Bel

[1] This seems to be reflected in representation systems for music, too, as surveyed by Wiggins et al. (1993).

[2] I am making here the perhaps unwise supposition that buildings do not *mean* anything, in any agreed sense of the word!

and Kippen (1992) and Ponsford et al. (FC) (among many others) indicates clearly that grammars of different kinds *can* successfully capture conventionally agreed musical structures, rhythmically and harmonically. No further argument, surely, is needed.

The title of Dempster's (1998) paper, "Is there even a grammar for music?", is interesting in itself, since it might be taken to suggest that the author is looking for some almost Platonic notion of an externally defined arbiter. This, however, is contrary to much linguistic and cognitive science thought: grammars are useful descriptive tools, but in using them (and their associated parsers and generators) the claim is not usually made that they *prescribe* the output of a linguistic system (*cf.* Sloboda 1998, page 23). A weaker (and more usefully and easily tenable) position than this is that they may – perhaps approximately – *describe* the structure of an organism's or a system's input or output, or its eventual internal state. Also, the title might suggest that Dempster is making the assumption of *universality*, which Meyer (1956) clearly refutes, in supposing that there might be "a [single?] grammar for music". So in these senses at least, Dempster seems to be arguing against a case noone has made, at least not recently.

It seems, then, that there are reasonable grounds for supposing that grammatical formalisms, borrowed from linguistics, may be useful for describing music, or, more precisely, for describing musical styles and conventions.

2.4 What exactly is a grammar?

A grammar is a finite set of rules which enables the symbolic description of the structure of a potentially infinite collection of structured symbols, and that is all. Grammars are useful because they can be concise and they sometimes help to shed light on the usage and function of such structure. They are *not* generally viewed as correctness tests for utterances in whatever language they describe; rather, they are used as aids in analysis (for whatever reason) of that language. Blacking (1984) puts this nicely into a sociological context:

> "Grammars are attempts to codify the regularities of structure that communities generate in order to give coherence to their communication and to enable individuals to share meanings."

3 Music and Semantics

3.1 What does "semantics" mean?

When discussing meaning, particularly in the context of linguistics, it is important to define one's terms. First, consider the word "semantics", which is often used interchangeably with "meaning". It has at least three different (but confusingly related) meanings, and before we can proceed, we need to say which we are discussing[3].

The first (and, arguably, original) technical meaning is attributed to Tarski, who used it to denote the mapping from the expressions of a logical theory to some representation of the "meaning", of those expressions, which are viewed as purely syntactic.

The second usage is in linguistics, and is the one closest to the common usage in the musical context. It is derived from the Tarskian logical idea. A linguist gives a semantics to her language by mapping it into some more or less logical formalism. These formalisms are often reducible to a standard formalism such as the First Order Predicate Calculus, so that they can be said to be understood in relation to a standard measure. The semantics of the linguist, then, is the syntax of the logician.

The third usage, not so relevant here, is by computer scientists, concerning the association of fragments of computer program with expressions in logic, in order to demonstrate and verify the behaviour of the program.

Both the logical and the linguistic versions of semantics rely on the ill-defined notion of "meaning". In particular, in cases where we have reference in a real or imagined world, we can say that, for example, the referent of "John" is a given person. This is not so easy when the domain of discourse is, for example, arithmetic – unless one is a Platonist. In any case, the assumption that meaning can be obtained by stating relations which are "about" something in some world is fundamental, and that brings with it another important point: stating a relation means "saying what is *true*". In this view of the world, the primary purpose of communication via language or logic is to make statements about what is true, and, *via* logical inference, to derive things which are true and unknown from things which are true and known.

How, then, are we to relate this notion of semantics, or meaning, based in reference to a world and in the notion of truth, to music? There is clearly a trivial level on which we can do so, in terms of programme music (*e.g.*, the deliberately bucolic stylistic references in Beethoven's 6th Symphony), but this does not really get us closer to any abstract meaning of music, not least because it rarely applies. More abstractly, the *leitmotif* might perhaps be construed as a form of quasi-linguistic reference,

[3] The definitions in this section are heavily influenced by an unpublished note by Prof. Alan Bundy.

though it is clear that this device was artificially introduced, and is not fundamental to musical experience. Even if we could reliably account for reference in this way, we are still left with the question of the attribution of truth values to musical structures. It really is not clear what would be the value of such activity.

A simple thought experiment will show the vagueness of the notion of "musical meaning". If I ask you to read a page of an unfamiliar novel and then ask you "What does that mean?", you can easily tell me. In particular, the request itself is easily understood as one of a small set of possibilities, involving more or less inference from what the text describes – witness the fact that I have been referring to "meaning" in the linguistic context for a page or so now, without problem. If, on the other hand, I ask you to listen to five minutes of unfamiliar music, and then ask you "What does *that* mean?", it is much harder for you to answer. I suggest that the reason is not just that it is hard to but music into words, but more that it is not immediately clear what I am asking.

I suggest, then, that "meaning" and, particularly, "semantics" are dangerous words to bandy about in the musical context, because they are already more or less formally defined in ways which are contradictory to the nature of musical understanding, and because the notion of "musical meaning" is dangerously ill-defined. Authors such as Meyer (1956) seem to suppose that the term is sufficiently well-understood that it can be used without acknowledgement or definition. I suggest that this is not so.

3.2 Connotation in Music

Notwithstanding the problem of terminology in "meaning" and "semantics", there is a very strong body of opinion that musical can indeed communicate something. The question is: what?

Meyer, in his seminal text, "Emotion and Meaning in Music" (1956), aims to explain the latter of the two subjects of this title in terms of the former. This is indeed the position which seems to be the most readily tenable in the context of emotional analysis: essentially, Meyer views the achievement of affect in a listener as the communication of meaning by music.

Krumhansl (1997) has shown that there are two ways in which music can achieve affect on a listener: by achieving direct emotional stimulus; and by the suggestion of emotional stimulus. In other words, a listener may react physically to music, showing various signs of emotional arousal, or she may be able to explain the feeling of the music, without showing signs of experiencing that feeling herself; of course, the two may co-occur. These responses come under the general heading of "affect"; what I am trying to pin down here is not the affect itself, but the property of the music which causes it.

On the other hand, Steedman (1996) has a very different notion of semantics[4], bound up with the expectation generated by particular chords in terms of a 3-D harmonic space.

I suggest, then, that we might be better off referring not to "musical semantics" or "musical meaning", but to *musical connotation.* Like any other existing word, "connotation" has a meaning already, but not formally in this context, nor in the other contexts in which "semantics" is used. Given the dichotomy of affect in Krumhansl's work, we might propose sub-terminology, such as "direct" and "indirect", to mean respectively the connotations directly inducing emotion, and those mediated consciously.

The questions then are: What is musical connotation? What causes it? and What is its relation to linguistic meaning or semantics, if any?

3.3 The Distinction between Syntax and Semantics

I return now to the assumption behind Dempster's (1998) question. What would a "grammar for music" be? In language, Chomsky (1965) tells us, the syntax of an expression forms a structure, on which meaning is superimposed, like a carrier wave supporting a signal, in radio technology. A syntactic category is then a distinct kind of entity from from a semantic one. But, while this approach can be useful for the purposes of studying language, it occludes a much more subtle view of the world, in which syntactic categories correspond precisely with semantic ones. For example, the *syntactic* category of "nouns" differs from that of "verbs" in a well-understood *semantic* way, simplistically stated as the distinction between objects and actions. Normally, linguists make a distinction between singular and plural nouns, which are indeed syntactically marked as distinct, by *sub-categorisation* of the "noun" category. This is a convenient device for keeping grammars simple. But in terms of meaning, the categories of singular and plural nouns are also clearly very different: there are many forms of reference which apply to one and not to the other. Similar examples of sub-categorisation abound, and in very many cases, one can see that it is semantics which drives the evolution of syntax.

Indeed, if this were not the case, then the very idea of a *compositional* semantics, so important to

[4] And it is to be noted that he wraps the word in scare-quotes!

computational linguists, would be a non-starter. The standard methodology is that one uses syntax to break up a string of symbols into a structural description; one then translates the symbols directly into individual representations of their meanings, and then *composes* them, according to the structural description, into a single statement of the combined meaning of the string. For this to work, one needs a semantics whose structure mirrors the syntactic one closely, as described above. As one builds more information (such as sub-categorisation) into one's syntactic description (and it is important to note that I am using the term "syntax" here to refer to the *categorisation* only – its expression in terms of word-level phenomena is strictly *morphology*[5]), the mapping between syntax and semantics moves from being one-to-many towards being one-to-one. This kind of thinking is apparent in recent moves in the computational linguistics world towards lexically orientated grammars – that is, grammars most of whose information is stored at word level, in the lexicon, rather than at an abstracted, distinct grammatical level. A good example of this kind of work is that of Steedman (1999), whose preferred formalism, *categorial grammar* (CG), I borrow for the following brief example.

In CG, syntactic categories are thought of as types, describing functions and constants. There are several ways to begin, but one standard is to start with nouns, **n**, noun phrases, **np**, and sentences, **s**. We can then write down the types of other categories in these terms. For example, a determiner (*e.g.,* the) is of category $\mathbf{np}/\mathbf{n}$, meaning that it combines with a noun to its right, to make a noun phrase. Similarly, a transitive verb (*e.g.,* drop) is of category $\mathbf{s}\backslash\mathbf{np}/\mathbf{np}$, so it combines with a noun phrase to its left and another to its right, to give a sentence.

The semantically important part is that these categories are associated with representations of constants and functions. A simple semantics might represent an object of type t as $o(t)$, and an action such as "x drops y" by the abstraction $\lambda y.\lambda x.drops(x, y)$, where λ marks variables to be filled in later.

So we might give the following lexicon:

man = **n**:*man*
ball = **n**:*ball*
the = $\mathbf{np}/\mathbf{n}$:$\lambda x.o(x)$
drops = $\mathbf{s}\backslash\mathbf{np}/\mathbf{np}$:$\lambda y.\lambda x.drops(x, y)$

and so, given the operation of β-reduction, which allows us to fill in the variables marked by λ, we can compose a semantic representation, as follows (the order of combination being determined by the category-types):

$$\text{The + man + drops + the + ball}$$
$$\downarrow$$
$$\mathbf{np}/\mathbf{n}{:}\lambda x.o(x) + \mathbf{n}{:}man + \mathbf{s}\backslash\mathbf{np}/\mathbf{np}{:}\lambda y.\lambda x.drops(x, y) + \mathbf{np}/\mathbf{n}{:}\lambda x.o(x) + \mathbf{n}{:}ball$$
$$\downarrow$$
$$\mathbf{np}{:}o(man) + \mathbf{s}\backslash\mathbf{np}/\mathbf{np}{:}\lambda y.\lambda x.drops(x, y) + \mathbf{np}{:}o(ball)$$
$$\downarrow$$
$$\mathbf{np}{:}o(man) + \mathbf{s}\backslash\mathbf{np}{:}\lambda x.drops(x, o(ball))$$
$$\downarrow$$
$$\mathbf{s}{:}drops(o(man), o(ball))$$

So now we have a representation in something very close to standard predicate logic, which which we can deal computationally.

This example demonstrates how the connection between syntax and semantics serves the computational linguist – the gross syntactic analysis (the category-types above) gives generous hints as to how to combine the representations of meaning to give the meaning of the whole. Of course, for a realistic grammar, things are much more complicated, but this is the basic idea.

To return to the analogy between linguistics and music: we can easily imagine how we might give structural descriptions of conventional musical forms, either in a way similar to CG, as above, or in a more familiar Chomskian style. This is mostly uncontroversial, and already has been well tested. But how can we posit an equivalent of the semantic element for music?

If semantics is something which achieves affect, as Meyer says, then perhaps we can find a representation which will allow us to denote it formally, as logic allows us to denote linguistic meaning; we need much more experimental data before we can begin such a task. If semantics is something which leads to expectation of subsequent harmonic structure, as Steedman cautiously proposes, then is it not really part of musical syntax? If the latter, then do we really need a distinct notion of semantics at all?

The point made above regarding the distinction between syntax and semantics is that, in language, syntactic and morphological structure mirror meaning at a much more subtle level than the broad syntactic categories within which linguists usually work. This leads to the suggestion that syntax and semantics, as discussed by computational linguists, should not be thought of as two separate (but related) entities – rather, they are actually different aspects of the same phenomenon. This view is supported by the success of lexicalised grammatical descriptions.

The question of the communication of musical

[5]This view begs the question: What determines word order? The answer is that syntax does, but sentence structure is a further aspect of syntax which is orthogonal to this part of my argument.

meaning has now become slightly less vexed. If syntax and semantics are viewed as (different aspects of) the same thing, then one can ask "how does this musical structure convey its connotation?" instead of "what does this music mean?". Another way to put this might be that music's meaning is *in* its structure, rather than being *carried by* its structure.

This leads us very comfortably back to Krumhansl's (1997) work, and that of others in similar vein. In a theory where we need an explicit "meaning", distinct from structure, there is a difficult gap to be bridged; from the point of view I espouse above, no such gap exists. Instead, we refute the supposition that "music must have meaning because language does", and admit a fusion of the (linguistic-style) syntax and semantics, allow the musical structure itself to generate affect directly (or indirectly, as Krumhansl points out).

4 Conclusion

In this paper, I have proposed an argument concerning the relationship between computational linguistics and music composition and analysis. I have suggested that the commonly assumed analogy between the linguistic syntax/semantics distinction and the idea of musical meaning is not a tenable one, and argued that a more simply structural approach may be more worthwhile. I have given a simple example to help demonstrate that syntax and semantics in language are not as distinct as some might suppose, and suggested that treating musical affect as a direct result of musical structure, rather than a phenomenon mediated by some further notion of meaning, is an appropriate way to proceed.

References

Angluin, D. (1980). Finding patterns common to a set of strings. *Journal of Computing Systems Science 21*, 46–62.

Baroni, M., R. Brunetti, L. Callegari, and C. Jacoboni (1984). A grammar for melody. relationships between melody and harmony. In M. Baroni and L. Callegari (Eds.), *Musical Grammars and Computer Analysis*, pp. 201–218. Firenze, Italy: Leo S. Olschki Editore.

Bel, B. and J. Kippen (1992). Bol processor grammars. In O. Laske, M. Balaban, and K. Ebcioglu (Eds.), *Understanding Music with AI – Perspectives on Music Cognition*, pp. 366–401. Cambridge, MA: MIT Press.

Blacking, J. (1984). What languages do musical grammars describe? In M. Baroni and L. Callegari (Eds.), *Musical Grammars and Computer Analysis*, pp. 201–218. Firenze, Italy: Leo S. Olschki Editore.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: The MIT Press.

Dempster, D. (1998). Is there even a grammar of music? *Musicæ Scientiæ 2*(1), 55–65.

Krumhansl, C. (1997). Musical tension: Cognitive, motional and emotional aspects. In *Proceedings of the 3rd Triennial ESCOM Conference*. European Society for the Cognitive Sciences of Music.

Lerdahl, F. and R. Jackendoff (1983). *A Generative Theory of Tonal Music*. Cambridge, MA.: The MIT Press.

Meyer, L. (1956). *Emotion and Meaning in Music*. University of Chicago Press.

Mitchell, W. (1992). The logic of architecture: programming the invention of physical facts. In *Proceedings of the 1992 Joint International Conference and Symposium on Logic Programming*, Cambridge, MA. The MIT Press.

Ponsford, D., G. A. Wiggins, and C. S. Mellish (FC). Statistical learning of harmonic movement. *Journal of New Music Research*. Forthcoming. Also available as Research Paper 874, from the Division of Informatics, School of Artificial Intelligence, University of Edinburgh.

Schenker, H. (1992). *Beethoven's ninth symphony : a portrayal of its musical content, with running commentary on performance and literature as well*. New Haven: Yale University Press.

Sloboda, J. (1998). Does music mean anything? *Musicæ Scientiæ 2*(1), 21–28.

Steedman, M. J. (1996). The blues and the abstract truth: Music and mental models. In *Mental Models In Cognitive Science*, pp. 305–318. Mahwah, NJ: Erlbaum.

Steedman, M. J. (1999). Categorial grammar. In *The MIT Encyclopedia of Cognitive Sciences*. MIT Press.

Wiggins, G., E. Miranda, A. Smaill, and M. Harris (1993). A framework for the evaluation of music representation systems. *Computer Music Journal 17*(3), 31–42. Machine Tongues series, number XVII; Also from Edinburgh as DAI Research Paper No. 658.