Chapter 1
Music Analysis by Computer: Ontology and Epistemology

Alan Marsden

Abstract This chapter examines questions of what is to be analysed in computational music analysis, what is to be produced, and how one can have confidence in the results. These are not new issues for music analysis, but their consequences are here considered explicitly from the perspective of computational analysis. Music analysis without computers is able to operate with multiple or even indistinct conceptions of the material to be analysed because it can use multiple references whose meanings shift from context to context. Computational analysis, by contrast, must operate with definite inputs and produce definite outputs. Computational analysts must therefore face the issues of error and approximation explicitly. While computational analysis must retain contact with music analysis as it is generally practised, I argue that the most promising approach for the development of computational analysis is not systems to mimic human analysis, but instead systems to answer specific music-analytical questions. The chapter concludes with several consequent recommendations for future directions in computational music analysis.

1.1 Introduction

The nature of music analysis, as a discipline, if not as a practice, has been a topic of debate on several occasions (e.g., Nattiez, 1990; Pople, 1994; Samson, 1999). Researchers in computational music analysis, on the other hand, tend to take music analysis as a kind of 'given'. My aim here is to revisit this ontological and epistemological debate, with two objectives: first, to draw conclusions useful for those who would use computers for analysis or who write analytical software; and second, to explore what insights follow from taking an explicitly computational approach to the debate.

Alan Marsden
Lancaster Institute for the Contemporary Arts, Lancaster University, Lancaster, UK
e-mail: a.marsden@lancaster.ac.uk

D. Meredith (ed.), *Computational Music Analysis*,
DOI 10.1007/978-3-319-25931-4_1

Fig. 1.1 Analysis as a mapping from piece to analysis

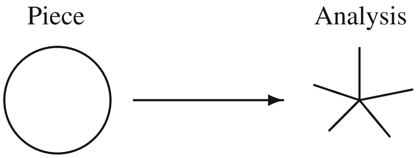

From a computational perspective, the simplest way of thinking of music analysis is as analogous to the working of a computer program. A piece of music is presented as input to a program, and this produces as output an analysis (Fig. 1.1). Music analysis, in this simple perspective, is effectively mapping from a piece to an analysis. As we will see below, this is too simplistic an account, but it will serve as a basis for the present. 'Mapping' here means there is a distinct relation between the piece and the analysis, involving also distinct relations between parts of the piece and parts of the analysis (see Sect. 1.4.1). It need not be a mapping in the mathematical sense, though it can often be expressed in that way. Different analyses of the same piece can exist because there can be different mappings, each corresponding to different programs. Any one program should always produce the same output for a given piece as input, unless it also takes input from some other, variable, source (such as a random-number generator). Human experts similarly map a piece to an analysis, and different experts produce different analyses. Different analyses arise through the application of different analytical approaches, whether these are gross differences between, say, analysis based on Schenkerian and Riemannian theories, or minor differences resulting from different interpretations of theory.

The following discussion will therefore be framed around three loci: the input or music, the output or analysis, and the mapping. Debates about analysis often go beyond discussion of the nature of music and analysis to also discuss the epistemology of the enterprise: how one comes to know what the mapping between music and analysis should be. In fact, although I say that the debate often goes beyond ontology, in writing about music analysis it often slips almost imperceptibly into epistemology. In what follows, I hope to make the move from one to the other more explicit, and also to draw out some definite conclusions and recommendations from the examination of ontological and epistemological questions.

1.2 Ontology of Pieces of Music

Questions of the ontology of music, of what it actually *is*, have generally focused on two issues. The first is whether music, or pieces of music, can properly be considered to have a distinct existence or not. At one extreme, a piece of music is considered to be a distinct 'object' with an abstract existence made manifest in various ways in the scores, recordings and performances of the piece, not to say also in the imagination of its composer and the memories of those who read the scores or hear the recordings and performances. At the other extreme, pieces of music are a kind of cultural fiction and music is ultimately an *activity* of which scores, recordings and performances are merely traces which exist only in certain cultural contexts (for discussion, see Goehr, 2007; Goodman, 1976; Levinson, 1980). Indeed, cultural context is of considerable importance in this debate, and it seems clear that different positions with respect to the nature of music, or at least the nature of pieces of music, are appropriate for different cultures and different periods of history. There are musical cultures in the world which have no scores, for example, and in which the notion of ‘a piece’, if it exists at all, is clearly something quite different from a symphony from nineteenth-century Europe.

The debate about the ontology of music, or of pieces of music, need not concern music analysts too deeply, though. Provided they show some due concern for the nature of the materials they use, they are able to proceed with their activities, and to produce useful insights, without having to commit themselves on the finer points of ontological debate. The computational music analyst, on the other hand, is not so free to use diffuse conceptions of a piece of music. The input to analytical software must be, ultimately, a binary code, in which every bit is unequivocally 0 or 1. Furthermore, every bit must be determined before the analysis begins, unless the analytical software is embedded in an interactive system which takes inputs from the user or other sources in the course of analysis, which is not typically the case for computational music analysis (see, however, Chap. 8, this volume, for an example of an interactive analysis system along these lines). The computational analyst is therefore effectively placed in a position at one extreme pole of the ontological debate, where pieces of music, or at least their manifestations used for analysis, are entities of which every feature is fixed. A human analyst can delay commitment about some features until part-way through the analysis. A curved line in a score, for example, might be interpreted as a slur or as a tie depending on the results of earlier stages of an analysis. One might object that the input to the analysis is then neither a slur nor a tie but a curved line, which is fixed in advance, but the example can then be moved down a level: is the curved line a proper part of the score or a printing error? There is no fixed boundary for the human analyst separating the information which can be used in the analysis from the information which cannot (see Sect. 1.2.1 below).

The upshot of this for computational analysis is that those who analyse pieces of music which clearly do *not* exist as fixed entities (e.g., from cultures without scores in which pieces are highly variable) should recognize that they do not analyse ‘the piece’, but rather some particular version or manifestation of it.

1.2.1 The Music Itself?

The classic contrast between music analysis and other areas of musicology used to be that analysis concerned itself with ‘the music itself’ rather than the music’s context: historical, social, etc. This changed in the late twentieth century with the advent of ‘new musicology’ which questioned the entire notion that music could be extracted from its context. Whether or not this is always the case need not concern us here, but it is clear that at least the boundary between what one must include in ‘the music itself’ and what one can leave out is far from clear. In some music, from Europe in the early 18th century for example, it is clear that performers would regularly add ornaments. In some cases, especially later in that century, ornaments are written in scores. When analysing a piece of music, should one take into account the ornaments or not? And if one takes into account the written ornaments from later in the century, why not the unwritten ones from earlier? If the ornaments are not written in the score, how can one be sure what they are?

Even once such issues have been settled, there is still uncertainty over what properly constitutes the ‘input’ to analysis. Here the issue is not so much what is in the music itself, but whether analysis really considers *just* the music itself. One would normally analyse a piece from the twentieth century differently from the way one would analyse a piece of music from the Renaissance, but the date of composition of a piece is not normally considered to be part of ‘the music’. Even if the focus of analysis is ‘the music itself’, it is clear that other information forms part of the input to the process.

Thus the boundaries for the input to the analytical process are indistinct in at least two ways:

- There is no clear definition of what information is included in ‘the music itself’ and what is excluded.
- It is not clear what additional information is required for proper analysis of a piece.

The simple diagram used in the introduction above should therefore be revised to show that the input to the analytical process, as traditionally conceived, has indistinct boundaries (Fig. 1.2). (Discussion of what a ‘proper’ analysis might be follows in later sections.)

As an example, consider the case of the ‘371 Bach chorales’, which were considered a kind of dataset of good practice in harmony long before ‘dataset’ became a word. It has been known for some time that (a) there are not 371 distinct chorales (some are duplicates, with varying degrees of similarity), (b) they are not all by J. S. Bach, (c) independent instrumental parts are missing from some, and (d) the precise details of some are uncertain (sources differ). It has nevertheless been common for computational analysts to use the version available in ‘Humdrum kern’ format (Sapp, 2014) without questioning its validity.[1] When a source is large in size

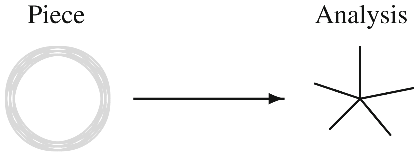

Fig. 1.2 Indistinct piece

[1] It must be acknowledged that the fault here is not Craig Sapp’s. He clearly states the source used in compiling the version (down to the precise print and edition) and states its nature as a collection made by C. P. E. Bach and Kirnberger.

and easily available, it tends to be used in several studies, presumably because the researchers want to put time into analysis rather than into encoding sources. For a discussion of the adverse consequences of unquestioning use of such sources, see Sturm (2014).

1.2.2 Differences and Indeterminacy in Digital Data

While computational analysis must, because it deals with digital representations, take a definitive input, this does not mean that it deals with single definitive representations of pieces of music. An audio file in WAV format does not contain the same stream of digital data as the same audio in FLAC format. Yet we are right to say it contains *the same* audio because exactly the same information is encoded in both files: a rendition of both into streams of digital samples would produce identical streams. Such differences between input data are merely differences of format and have an effect only at an operational level; there is no reason to expect any difference in the analyses produced. We might consider the two representations to be two different 'projections' of the same data by analogy with the projection of graphical data from one co-ordinate space to another: the data is not changed but it is represented differently.

Other differences in input to computational analyses are more significant, and even when the word 'format' is used in connection with these, the differences constitute genuine differences in the information represented. An audio recording in a lossy compressed format such as MP3 is different from a recording in a non-compressed format because some of the information in the latter has been lost. A representation of a piece in a MIDI file and a representation of the same piece in a MusicXML file contain different information. For example, the MIDI file does not contain information about the location of barlines relative to the notes, whereas this information is essential in the MusicXML file. Differences of this kind are not so much ones of projection but ones of *selection* of data. Choosing to use one 'format' or another as input to the analytical process is tantamount to selecting some information about the piece and ignoring other information. Of course the data has to be present in the source which is being represented in order to be selected, so while one can take a MIDI file direct from the playing of a musician on a keyboard, one cannot directly derive a MusicXML file from that source. The location of barlines is not explicit in the musician's playing on the keyboard, but it is explicit in the score, and one can derive a MusicXML file from a score.

This difference with respect to sources underlines another kind of difference in different 'formats'. This is clearest with respect to the representation of pitch. A score shows the pitch of a note in a particular spelling (e.g., C♯ or D♭), and formats such as MusicXML contain similar information. A MIDI file, by contrast, contains pitch information only in relation to which key on a keyboard is pressed (so C♯ and D♭ are represented as the same pitch). On the other hand, a MIDI file contains information about precise timings, whereas a MusicXML file represents timings only by the notated durations of the notes and a generally imprecise indication of tempo. One can unequivocally derive the pitch information necessary for a MIDI file from a score, but one cannot derive the precise timing information required (unless the score indicates a precise tempo) without making a guess at an appropriate tempo, and using that as a basis for timings in the MIDI representation. Such an interpretative step is often required even when creating a MusicXML representation from a score, despite the fact that MusicXML was created precisely to represent score information (Good, 2001). When simultaneous notes occur, the format requires a distinction to be made between chords and simultaneous notes in different voices. Different people are likely to make different decisions about which is appropriate in some cases. We might use a single notated source as a guide to aim at definitive decisions, representing simultaneous notes as chords when they are attached to the same stem but in different voices when they are not. This does not overcome the problem of ambiguities in the assignment of notes to voices, however, and most obviously does not help in cases of simultaneous notes without stems such as semibreves (whole notes). Ideally, there would be conventions for making such decisions in a regular fashion, and so allowing the formation of canonical MusicXML files, but I am not aware of any such conventions.

In the cases of some computational music data, there is no evident interpreting agent. When scanning a score to generate a digital image, for example, the user places the score on the scanner and presses a button rather than taking a set of decisions on how to represent the information in the score. However, the resulting digital data is still contingent on factors which, like the decisions of an interpreting agent, are specific to the manner in which the scan is made rather than dependent solely on the score itself. The score might be placed at an angle, for example, so that the staff lines do not appear horizontal in the image. In the case of audio recordings, it is well recognized that factors like the placement of microphones and the acoustics of the recording space have a profound effect on the resulting data. Microphones and scanners, and indeed any sensing device, have an inevitable element of noise in their outputs, constituting another source of information effectively added to the original music data. Often one can go some way towards removing such unwanted data (often called noise or distortion)—a scan can be de-skewed, for example, or noise-reduction applied to a recording—but to do so relies on assumptions about the original data (e.g., that the staff lines are horizontal). There is no generalized way for distinguishing between original and introduced data on the basis of the data alone, and, perhaps more importantly, no way of distinguishing between introduced data which is possibly legitimate interpretation and introduced data which is spurious. In a recording, for example, how could one know if the sound of a particular instrument is prominent because the recording engineer has legitimately placed the microphone to best pick up the lead instrument, or because the resonance of the recording space happened to make it prominent?

There are thus three ways in which computational representations which form the input to computational analysis, while fixed in the sense that they have a definitive digital form, are nevertheless indefinite in the sense that there is not a single definitive form for the data representing a single piece:

- the projection of the data to a particular format,
- the selection of the data which is represented in that format, and
- some unavoidable element of extrinsically introduced information, or at least uncertainty about whether or not there is extrinsically introduced information.

The diagram of music analysis therefore needs revising once again, to reflect the varieties of possible inputs, even for a single, indistinct, piece of music. Unlike the indeterminacy discussed in Sect. 1.2.1 above, these are inputs which have distinct boundaries, but the indeterminacy arises from the impossibility of determining which is the proper input (Fig. 1.3).

In practical terms, the indeterminacy might be very low. We might, for example, decide to represent only the pitch and timing information from a score which we consider to be unequivocally given by the unambiguous placing of the notes vertically on the stave and horizontally in relation to barlines and each other. Of course we have selected information from the score, and we have implicitly assumed that the ignored information (articulation, for example) is not relevant to the analytical distinctions we aim to make, but at least we have not introduced anything which is not in the score. Or so we suppose, because while we might be confident that nothing has been introduced, we cannot be certain. Mistakes happen, and if we are using data generated by someone else, what information do we have about how likely it is to contain mistakes? More importantly, we cannot be certain that what is unequivocal in one musical source is unequivocal in another, so we cannot be confident that our procedure will generalize to all cases. For example, Don Byrd has given examples of standard music notation (by well known composers) where the timing of notes is ambiguous and some in which even the pitch is questionable (Byrd, 1984, 2013).

1.2.3 Error and Approximation

The solution to these issues is for computational music analysts to recognize that they operate with approximations. In the natural sciences, researchers are well used to dealing with input which contains noise and error. Every measurement is regarded as an approximation of the real value. Extremely reliable results can nevertheless be derived because researchers quantify the error in their measurements, take multiple measurements, and take these into account in calculating their conclusions. In the social sciences, too, the idea of multiple measurements to increase accuracy is common. To take an opinion poll, one asks not just one person their view, but many people, and one is careful to ask a variety of people so that the answers are not biased

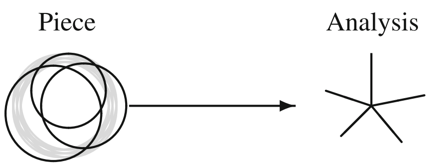

Fig. 1.3 Selected but definite input for analysis

to one point of view. The essential idea is that the overall properties of the *sample* one uses should match, within an acceptable degree of error, the overall properties one would find in the entire *population* if one were to actually measure everyone or everything in the domain.

In music analysis, one is generally concerned with the properties of a particular piece (sometimes even of a particular performance) and so some of the most obvious analogies of sampling from a population do not apply. However, a useful strategy might still be to consider different editions of a piece of music as sources, and perhaps performances as well as scores. It would be possible, too, to perform similar analyses on different formats of data, but since analytical software is usually specific to a particular format, and since it is not clear how 'the same' analysis can be ensured given data in different formats, this too is problematic.

The notion of error, however, is eminently applicable. In Fig. 1.3, error corresponds to the region within the fuzzy boundary of 'the piece' but outside the definite boundary of the particular digital representation used as input to the analytical process. Just as a natural scientist aims to minimize error, the computational analyst should aim to minimize the area of this region of the diagram. Of course, this cannot be done with certainty because the boundary of 'the piece' is indistinct. Furthermore, there is no definitive way of even estimating the area of the region. How much 'space' in the input domain is taken up by the pitch-spelling information which is not included in a MIDI representation, for example? The answer, if there is an answer, depends on the one hand on the nature of analysis undertaken, and on the other on the viewpoint, implicit or explicit, of the person judging its validity. A MusicXML file gives no information about the layout of a score on the page, and in many cases it is not regarded as significant: the area, in the sense of 'quantity of significant information' in the piece-as-score, which corresponds to the layout information missing from the MusicXML file will be close to zero. If the piece in question is *Eight Songs for a Mad King* by Peter Maxwell Davies, in which the staves of the third movement are arranged on the page in the shape of a birdcage, the area for information corresponding to layout will be greater.

A more effective way of quantifying error is to consider the effect on the analysis produced. How different are the analyses arising from different selections or interpretations of data from the piece? As indicated above, actually using different input formats will not always be realistic. However, one could simulate these differences to some degree. For example, as stated above, a MusicXML file, or indeed any input derived from a score, does not contain precise timing information. (Even if the score gives a metronome marking for the tempo, this is generally regarded as a suggestion for the tempo rather than a strict instruction to the performer.) When using analytical software such as Melisma (Sleator and Temperley, 2003) which requires timing information for input, it would be easy to make multiple runs of the software with variations in the timings in the input, made with the aim of producing a distribution of different timings similar to the distribution occurring in different likely interpretations of the information in the score. Comparison of the resultant outputs would then give good data on which to quantify the significance of error in the input. If the difference in the outputs is small, the error in the input can be considered to be small also.[2]

In some other cases, such as the Bach chorales mentioned above (Sect. 1.2.1), one can have some confidence in the smallness of the error even without this kind of testing. One could argue that (a) the number of duplicates is small in comparison to the number of chorales (371), (b) the number of known mis-attributions to J. S. Bach is tiny in comparison to the number of cases where J. S. Bach is confirmed as the composer, (c) very few chorales have missing instrumental parts which differ from the vocal parts, and (d) there are few differences between sources. However, even this kind of reasoning is usually absent from computational analyses, and researchers have tended to perform analyses as if there were no error in the input. The consequences of this oversight in the domain of genre recognition have been set out by Bob Sturm (2014). While researchers have claimed that their computational systems recognize the genre of audio recordings with a high degree of accuracy, an examination of the errors and inconsistencies in the dataset used and the effect of these on the measures of accuracy led Sturm to conclude that these claims are not supportable.

1.3 Ontology of Analysis as a Kind of Writing

The question of the ontology of an analysis is in some ways simpler and in some more complex than the ontology of a piece of music. It is simpler because the form is more distinct: analyses are pieces of writing or some other kind of discourse. They are communicated from one human being to another. (Or at least they were until computational analysis came along. It is entirely possible for the object of computational analysis to involve no human reading of the analytical outputs, but discussion of this will be left for later. For the present, it will be assumed that the objectives of computational analysis are the same as those of analysis without computers.)

The question of the ontology of an analysis is more complex because analyses are not surrounded by the same richness of data from patterns of use and context as are pieces of music. We have a clear idea of the situations in which music is typically created and listened to, we have information about what music is valued in which situations, and we have a large body of commentary to draw on. For music analyses, we know that these are published in books and journals, and discussed in conferences and classrooms, but clear information on the role of this activity within wider culture is not clear. Analysis has regularly been used as a tool in the teaching of composition, and analysts claim that their work can inform performance of a piece (Mawer, 1999), but there is precious little evidence of such influence in everyday practice.

[2] David Meredith followed a procedure somewhat like this to test the impact of tempo on pitch-spelling using Melisma, finding the effect to be quite large (Meredith, 2007, p. 205). This was over a range larger than the tempo variations one would typically find in performance, however.

1.3.1 Description and Explanation

Analysis is typically distinguished from other kinds of writing about pieces of music—commentary, criticism, hermeneutic exegesis, etc.—by claiming that it considers the piece itself (which as we have seen above is not entirely true), that it eschews the value judgements inherent in criticism, that it avoids the subjective perspective of hermeneutics, and that above all it provides not just a description of a piece of music but an *explanation*. The distinction between description and explanation, however, is not so easy to make, and might be only one of degree: if we can describe something very succinctly, then we can thereby, in a sense, explain it. (Applications of this idea using minimum-description-length coding, information theory and Kolmogorov complexity will be briefly reviewed in Sect. 1.5.1.)

In the domain of the natural sciences, a phenomenon is explained by demonstrating how general principles apply in a specific case. The characteristics of a species, for example, might be explained by demonstrating how the principles of the theory of evolution apply to the species in its particular environment. This kind of thing exists in music analysis also, where a piece is explained through a demonstration that it follows a particular model, e.g., Sonata Form. Herein lies a paradox, though. Music analysis is generally distinguished from music theory, and the distinction is that analysis is concerned with particular pieces of music whereas theory is concerned with generalities. An analysis, from this perspective, seeks not just to explain a piece by reference to a general model, but by reference also to *its particular properties*. While apparently seeking to *explain*, then, an analysis of a piece of music often takes a great deal of space to *describe* the particular and distinctive characteristics of that piece. Although I stated above that analysis eschews value judgement, value is rarely far from the surface. The apparent objective of many analyses (and in some cases, such as in some of the analyses of Schenker, the explicit objective) is to demonstrate how a piece of music is a masterwork. To do so requires pointing out its special characteristics, not just the ways in which it follows general models.

One response to this is for analysis to seek how a piece establishes its own explanatory principles, or it might seek to build an explanation either on scrutiny of listening experiences (one's own or others') or on information concerning the piece's creation. Nattiez famously distinguishes three levels for analysis: the poietic, which involves consideration of the process of creation; the esthesic, which involves consideration of the process of hearing; and between them a neutral level (Nattiez, 1990). Analysis at the neutral level examines the divisions and patterns in the neutral code, which in practice for most music analysis means the notes written in the score. An important component of the method of 'paradigmatic analysis' advocated by Nattiez is the discovery of 'paradigms' within the piece. These paradigms are, at first, based on evident similarities in the configuration of notes; but later, they take on a more generative role in determining which configurations count as occurrences of a particular pattern and which do not (for a large-scale example, see Nattiez, 1982). The basis for organization of the piece is thus established *in the course of analysis*. On the other hand, one does not start an analysis (and perhaps cannot start) from a blank, theory-free position and follow only the leadings of the information at the neutral level. Certain principles, usually unstated, govern the establishment of paradigms in the course of the analysis. The same thing can be seen in a simpler fashion in the manner of motivic analysis as practised by Réti, Keller, and others (Keller, 1965, 1985; Réti, 1962, 1967). They start from an underlying principle that pieces of music are organized by the use of melodic material which derives from a single basic motive. The nature of the motive and the methods of derivation will vary from piece to piece, and ad hoc arguments are presented in the course of the analysis to justify the analyst's decisions. Essentially, motivic analysis and paradigmatic analysis operate with a meta-theory which governs the generation, in the course of analysis, of a specific theory to explain the piece of music in question. Those familiar with machine learning, genetic programming and other inductive systems might see a similarity here: an overarching principle (the minimization of a particular error function, for example) is used to guide the development of the analytical process, the details of which are contingent on the properties of the actual data. (Other points of contact between machine learning and computational analysis will be explored further in Sect. 1.5.2.) In the case of motivic and paradigmatic analysis, however, the meta-theory is usually only vaguely expressed at the outset and is subject to revision in the course of making the analysis.

1.3.2 Limitations of Mechanistic Analysis

Nattiez (1990, p. 32), though, claims that analysis at the neutral level is descriptive, and that it is poietic and esthesic analyses which are explicative. Exactly why this should be so is not clear. One possibility is that explanation requires reference to wider realms of human meaning, and this is not possible without consideration of either the poietic or esthesic level. Another possibility (though probably not one Nattiez would endorse, for reasons set out below) is that analysis at the neutral level is mechanistic, and explanation by a mechanism is not possible, perhaps again because explanation requires some human reference.

Nattiez does consider the possibility of mechanistic analysis, but in the earlier article he goes beyond the assignation of mere description to call the result of a mechanistic operation an *inventory*, which he contrasts with an analysis. For Nattiez, analysis requires a step of deciding which relationships between musical units are to be considered transformations, placing the musical units into the same 'paradigm', and which relationships are to be considered distant enough to distinguish one paradigm from another.

> Given that not all possible forms of transformation are foreseeable, as soon as relationships are established between units that are not strictly identical we enter the realms of analysis. [...] The difference between an inventory and an actual analysis is that *it does not appear to be possible to deduce the latter from the sum of the information provided by the former*.
>
> (Nattiez, 1982, p. 256)

This is quite a strong challenge to computational analysis. Some will take heart from the fact that Nattiez's claim was made when computational analysis was in its infancy, and that Nattiez did not have a clear idea of what it was possible to do with a computer. Decisions on whether relationships count as class membership or not are now made many, many times a day by search engines. However, there is one further aspect of Nattiez's argument which will be readily recognizable to computer scientists. Later on in the same article, Nattiez acknowledges that mechanistic accounting can treat not just identity but any kind of relationship between musical units, expanding the inventory to accommodate these by adding extra 'columns' in a table to record the occurrence of particular units or transformed units. Extending the point quoted above that "not all possible forms of transformation are foreseeable", Nattiez points out that "[t]here is, therefore, no limit to the number of possible columns" (Nattiez, 1982, p. 257). Since every real computing machine has finite resources, we therefore can never be certain that the relationships required to make a proper analysis of the piece will be in the table.

There are, however, two flaws in this argument.[3] First, it is only true that the table of possible transformations will be infinite if the set of units to be related is itself infinite. If the representation of the piece at the neutral level is finite (e.g., a finite sequence of notes described in terms of pitch and duration), then the set of all possible relations between notes or sets of notes is finite (though large).[4] It would be possible for a computer with sufficient resources to make an inventory of all these relations and, given a mechanism for deciding which relations were of significance, to derive an analysis from this inventory, in Nattiez's terms. In most cases, however, the inventory would be impossibly large for this to be a realistic method of analysis.

The second flaw is not one which renders the argument invalid, but one which blunts its force. It was argued above (Sect. 1.2.3) that the fact that one cannot be certain that the input to computational analysis covers all the necessary information about the piece does not render analysis impossible. By a similar argument, the fact that one cannot be certain that all significant relations can be recognized by the analysing computer does not render analysis impossible. Once again it is a matter of statistics. How confident can one be that the significant relations can be taken account of? The fact that the same kinds of relationship seem to be regularly reported in music analyses suggests that it is possible to design analytical software a priori which is likely to encompass the significant relations. To extend this argument to quantify the degree of confidence would require some idea of what makes a relationship significant, which is moving towards questions of epistemology to be discussed below. As with the inputs, though, it would be possible to estimate confidence in the analysis on the basis of experiments using different sets of relationships in the early stages of analysis (the 'inventory' stage in Nattiez's terms). If the analysis

[3] It has to be acknowledged that Nattiez adumbrates rather than explicitly states this argument, so my interpretation of his meaning might be incorrect.

[4] Note that a relation in this sense is defined purely by the set of sets of notes which show that relation. There does not need to be a definition of the relation in terms which would allow us to determine, for any arbitrary set of notes which do not actually occur in the piece, whether they are in that relation or not. For example, if in a particular piece there is a relation consisting of the two pairs of notes with pitches (A, B) and (B, A), there is no necessity to define whether this relation is defined as 'transpose the first note up one step and the second down one step' or 'swap the first and second notes', or indeed in any other possible way.

uses pitch intervals as a basis for segmenting a sequence of notes, for example, one could experiment with runs of the software which use absolute intervals expressed in semitones or key-dependent intervals expressed in scale steps and investigate how much variation there is in the resulting segmentation.[5]

1.4 Ontology of Analysis: What Is It About?

An analysis is not adequately defined as a piece of writing; we want crucially to know what it is writing *about* when it describes or explains, and what kinds of things it says. Just as there are multiple answers to the question of the ontological nature of a piece of music, there are multiple perspectives on the ontology of an analysis.

1.4.1 Temporal Basis of Analysis

Before embarking on discussion of ontology proper, it is worth clarifying certain general characteristics of an analysis. As discussed above, an analysis contains information derived from a piece of music. This is not typically a single piece of information, or a piece of information which applies to the entire piece. There are other kinds of derivation of information from pieces which are not analyses, as for example determining who wrote a piece where there are disputed attributions, or determining the genre of a piece of music. Analysis can contribute to making this kind of determination, but the determination itself is not analysis. Even to say something like ‘piece X is in Sonata Form’ does not constitute an analysis, though it is a kind of analytical statement. The analysis would show us *how* the piece is in Sonata Form.

The distinguishing characteristic of an analysis is that, like the piece of music, it has a temporal structure (or at least a structure which maps onto time), and the mapping shown in Figs. 1.1–1.3 is not just from the piece to the analysis but from parts of the piece to parts of the analysis, and the temporal relations between the parts of the piece are reflected in some way in the relations between the parts of the analysis.

1.4.2 Information Content of Analysis

A second characteristic of analyses is that they contain, in a technical sense, less information than the piece analysed. An analysis does add information in the sense

[5] This resembles the approach of testing the effect of selecting different combinations of multiple viewpoints (Whorley et al., 2013) or of parameters (van Kranenburg, 2008), but generally this has been done in order to optimize the accuracy or efficiency of the result rather than in order to estimate its degree of error.

that someone reading an analysis will gain knowledge about the piece which they did not have before, but this is a gain for the reader brought about by directing his or her attention to aspects of the piece or derivations made from the information in the piece. Nothing is added which was not already latent in the piece, and a lot is left out because many alternative ways of structuring the piece have been excluded. On the other hand, whether one regards an analysis as always containing less information than the piece depends on one's view about what information is contained in 'the piece' in the first place and requires the questions about the ontology of a piece of music from Sect. 1.2 to be revisited.

As mentioned above (Sect. 1.2.2), the input to an analysis is already a *selection* from the information in the piece, but, as was also indicated above, analysis generally goes beyond this to reduce the quantity of information further. This is most obvious in a Schenkerian analysis, where the term 'reduction' is used explicitly, but it is evident in other kinds of analysis also, such as those which explain a piece in terms of a small set of melodic motives. Indeed, it can be argued that a necessary consequence of giving an explanation of the piece rather than merely a description will involve reduction in information. The discussion of significant relations in Sect. 1.3.2 is of relevance also: selection of significant relations will entail a reduction in information. This topic will be revisited in the discussion of epistemology below.

1.4.3 Analysis as Inherent Within the Piece

Analysis is often presented by its practitioners as demonstrating how a piece of music 'works' (Bent, 1987). Analogies are sometimes made to architecture (in phrases such as the 'plan' or 'structure' of a piece), to anatomy (the 'skeleton' of a piece) or to growth (the 'germ' of a piece). The analyst is like a surgeon dissecting a body to reveal its skeleton, ligaments and organs. Indeed, some find music analysis distasteful on similar grounds to dissection: to them the uncaring cutting and poking of a living thing, or at least a once-living thing, is a kind of violence. The common riposte is that, unlike dissection, the act of analysis does no damage to the piece which can, as it were, arise from the operating table and walk again. If music analysis does do damage to a piece, to my mind that is because some erroneously believe it prescribes a way of hearing. There is an important distinction between the act of analysis and the act of listening.

In the perspective of analysis as revealing structure, the analysis reveals what is already inherent in a piece, in the same way as a skeleton is inherent in a body. This follows also from the perspective of analysis as demonstrating how a piece 'works', which is not inconsistent with different analyses of a piece. We might encounter different descriptions of how a piece of machinery works, for example explanation of a hydraulic jack in terms of the incompressibility of the fluid and the volumes displaced or in terms of the conservation of energy, and we can see the correctness of both while still believing that there is *one* way in which the machine works. The different explanations offer different ways of accounting for how the machine works.

An analysis, in this perspective, has a relation to 'the piece' similar to scores or performances of the piece. There is a single entity which is 'the piece', and multiple manifestations which serve particular purposes, selecting information inherent in 'the piece' and projecting it to be presented in a particular way. These acts of selection and projection are themselves determined by other information, such as the date of composition of a piece, or a theory of tonal structure, but the information which the analysis presents is, essentially, information inherent in 'the piece' (Fig. 1.4).

1.4.4 Analysis and Cognitive Structures

Just as the objective existence of a piece of music is questionable, so too is the existence of a non-contingent analysis. The analysis itself might be a kind of document, which can have an existence as objective as a score, but the 'analysis' which this document conveys, the significant information, exists, perhaps, only in minds (Fig. 1.5).

There is a kind of analysis which tries to explain a piece by reference to the real or supposed processes and structures in the mind of its composer. This is Nattiez's poietic analysis, and we find analyses which take as supporting information such documents as composers' sketches. In recent decades we have set less store by the authority of composers, and so analysis which relates explicitly to cognitive structures more commonly now refers to the minds of listeners. This is like Nattiez's esthesic analysis, but the referents, at least for computational analysts, are generally not semiological but psychological. The approach of Lerdahl and Jackendoff, for example, explicitly aims to examine "the musical intuitions of a listener who is experienced in a musical idiom" (Lerdahl and Jackendoff, 1983, p. 1).

1.4.5 Analysis as Interpretation

One danger with the conception of analysis as revealing something of the cognitive structures engendered in listening to a piece of music is that it assumes that there is something fixed to reveal. The boundaries of 'the piece' in the listener's mind are likely to be even more fuzzy than those of an objective 'piece', and we have to contend also with multiple 'pieces' in the minds of multiple listeners.

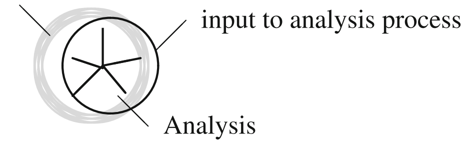

Fig. 1.4 Analysis reveals structure inherent in the piece

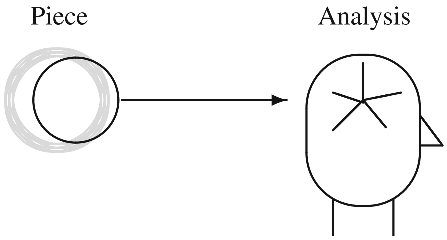

Fig. 1.5 Analysis as cognitive structure

Several commentators, therefore, have abandoned any idea of analysis presenting an authoritative explanation of how a piece ‘works’, and replaced this with either a subjective ‘works for me’ or a shared or inter-subjective ‘could work’. The analysis is explicitly not *the* analysis of the piece, but one possible analysis. If we retain a stance that the analysis shows structure in the piece, then we must abandon the analogies of architecture and anatomy. The structure is no longer something definite in the piece, but something which is *constructed* by the process of analysis. The relation to listening persists, because listening too, in this perspective, constructs a structure, and the structure in the analysis might match this one. A relation with the structure latent or inherent in ‘the piece’ is not abandoned either: the configurations of notes and sounds which convey the music to the listener, and which are the raw material for the analyst, enable certain structures to be constructed, and hinder others.

One of the most eloquent advocates of this approach to analysis is Jim Samson, who states that the reduction from ‘surface’ to ‘structure’ (the analytical process under discussion here)

> far from providing an empirical explanation of a work, can only offer an interpretation of it, albeit one which may be constrained by something akin to a rule-governed system.
>
> (Samson, 1999, p. 45)

The resulting analysis depends on

> the role of observer, who [. . .] creates a theoretical predetermined and pre-analytic concept of the object to be analysed.
>
> (Samson, 1999, p. 46)

Samson goes on to relate this to the loss of faith within the natural sciences in ‘the stability of object description’. Some might therefore justifiably regard the difference between theory-based ‘explanation’ offered by the natural sciences and theory-based ‘interpretation’ offered by music analysis to be just one of degree. There are several important contrasts, however, that suggest that the difference lies not just in the degree of impact of the observer’s role, but also in a real divergence in method and focus.

1. The natural sciences generally examine many instances of a phenomenon; music analysis typically examines a single item.
2. The natural sciences derive principles to be applied in an objective and formulaic fashion; music analysis always presumes some degree of prior experience and understanding in the analyst and reader.

3. The principles of the natural sciences are tested by experiment in the real world; the principles of music analysis are generally tested against the agreement of other musicians.
4. Explanations in the natural sciences are acknowledged as imperfect, but the imperfections, embodied in differences between the outcomes of different plausible explanations, are minimized to a level where they can be safely ignored (to use Newtonian mechanics in the design of an aircraft is acceptable, for example, even though it is known that quantum mechanics is a more accurate explanation of physical phenomena); in music analysis, the differences between the outcomes of different plausible explanations are often the focus of debate, whose objective is not revision of principles to minimize error but rather to see one explanation prevail over another or, in more co-operative forums, to explore the different aspects of a piece explained by different analyses.

In music analysis, the analyst and reader are, or perhaps should be, always aware of what is *not* said, how the analysis might have been different, a point made strongly by Jonathan Dunsby in the editorial of the inaugural edition of *Music Analysis* (Dunsby, 1982).

In the conception of music analysis as interpretation, the structures shown are not simply 'discovered' in the piece but constructed, and the input to the process of analysis which is *not* part of the piece becomes crucial. (Note that, in practice, this input may not be recognized strictly as 'input' to the analytical process but may instead be embodied in the adaptation or design of the process itself prior to or in the course of analysis. It is nevertheless a kind of input, analogous in computation to the setting or modification of parameters.) Other perspectives on analysis do not deny that factors which are extrinsic to the piece are involved (as discussed above in Sect. 1.2.1), but these factors are regarded variously as a definite context or as the 'proper' environment of the piece. In the perspective of interpretation, the extrinsic factors are explicitly variable, and particular to the analysis: 'this is the way *I* hear this piece', or 'if we compare this piece to X (some other piece, or some theory of structure), the following pattern emerges'. To display this conception of analysis in graphical form, an additional arrow is required to show the crucial additional input to the analytical process, input which generally comes from the analyst herself or himself (Fig. 1.6).

I suspect most analysts follow this interpretative conception of analysis, even when they write as if they are revealing a structure intrinsic to the piece. (Perhaps

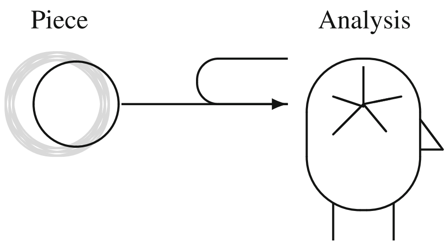

Fig. 1.6 Analysis as interpretation

they write in such a tone partly to imbue their writing with an aura of authority rather than because they believe other analyses are wrong.) Certainly this is how most analyses are used by their readers. Only students, I suspect, read an analysis to find the 'right' answers about the structure of a piece. Other analysts continually test their own conception of the piece's structure against what they read (at least, that is how I read analyses), and performing musicians seem to approach analyses with the same kind of curious but provisional approach taken to the performances of other musicians. Effectively, the analysis acts not to inform its readers of the piece's structure, but to propose a structure with which the readers may concur, often going so far as to seek to persuade readers to concur.

1.4.6 Computational Perspectives

Computational analysis, at first sight, seems most obviously to follow an ontology of analysis revealing the structure within a piece (Fig. 1.4). The input is data taken from the piece, and the analytical process generally finds specific structures (e.g., harmonies) or patterns (e.g., recurrences) within that data. Provided we accept the body of data as in some sense representing the piece, there is no reason to place the analysis in a locus different from the piece itself. On the other hand, there is also not necessarily any reason to believe that the particular structure in the output of the analytical process has a privileged status among the myriad other possible ways in which the input data could be structured. In other words, what authority does the computational process have to claim that the output is *the* structure of the piece or even that it is a significant, interesting or useful way of structuring the piece? (For some possible answers to the second half of this question, see Sect. 1.5.1.)

In the case of analysis by a human expert, the claim for a privileged, if not authoritative, status for the structure shown is implicit in the human's expertise. We tend not to ascribe expertise to computers, especially not in matters to do with the arts, so we are unlikely to regard computational analyses as carrying any particular authority. On the other hand, it is absolutely clear that the analysis which is the output of a computational analytical process *does* have a particular and unique status in the universe of possible analyses of the piece: it is the output of that particular analytical program when given that particular piece as input. If we know we want to find the structure which follows from a particular process, then a computational analysis, because it is free from bias, is *more* authoritative than a human one (provided the software properly implements the process). (In this regard, see also my analogy of computational analysis and forensic science in Sect. 1.5.2 below.)

Computational analysis can also sit comfortably with an ontology of analysis as revealing the structures of human musical cognition, though the epistemological issues around the correctness of the analysis are different from those arising from an ontology of analysis as revealing the structure in the piece (Sect. 1.5.1). The ontology with which computational analysis seems to sit least comfortably is that of analysis as interpretation. As discussed in Sect. 1.4.5, this kind of analysis places emphasis on the inputs to the process which are extrinsic to the piece and which cause one analysis to be different from another, even when the two share common music-theoretical bases. In the case of human expert analysis, these inputs come from multiple sources: the analyst's experience of listening to the piece, knowledge of other pieces and analyses, etc. To include all such possible inputs in a process of computational analysis would require modelling the entire human expert in all his or her aspects.

1.5 Epistemology

Some slippage from ontology to epistemology, which I remarked at the outset was a characteristic of writing on music analysis, can be detected also in the paragraphs above. From consideration of what an analysis *is*, one naturally and perhaps unavoidably moves to what it contains and why it contains what it does: from considering the nature of knowledge constituted by an analysis, one moves to considering the bases of that knowledge. The epistemology of music analysis in general is extremely fraught, so I will concentrate on some aspects of the epistemology of computational musical analysis only.

1.5.1 Correctness of Analysis

How do we know when a computational procedure produces a correct analysis? In one sense the analysis is always correct (provided the computer has not malfunctioned). That particular analytical procedure is not capable of producing any other output given the same input (remembering, in the case of stochastic procedures, that a random-number generator constitutes a kind of input). We are hardly likely, however, to set great store by the analysis of a piece produced by software which finds the structure of a melody in an entirely arbitrary manner such as to take every third note of the melody. The structures we want to find, as discussed above, have some special status which might be determined on one of several different bases. A more creditable analysis is likely to result from a procedure with a more sophisticated basis, and especially one where there are good grounds for believing that the basis implies musical significance also. A particularly promising approach involves a group of concepts around compression, minimum description length, information theory, and Kolmogorov complexity. As mentioned above (Sect. 1.3.1), there are good grounds for taking a parsimonious description of any phenomenon, including a piece of music, to be an explanation of it, and the more parsimonious the description the better the explanation. Software can find the description or 'analysis' of a piece which is the most parsimonious or at least close to the most parsimonious, and is therefore, in those theoretical terms, 'correct'. For applications of these ideas in computational music-analytical research, see Pearce and Wiggins (2012), Meredith (2012) and Chaps. 7 and 13 in this volume.

One way we might know if a computational analysis is correct is if its output matches analyses produced by human analysts. This is a path taken by several projects in computational analysis, for example (Anglade et al., 2010; Marsden, 2010; McVicar et al., 2011; Pauwels and Martens, 2014). In such cases, though, nothing is added by the computational analysis to explanation of the pieces analysed. An analytical dividend would come only from application of the same computational procedures to other pieces, and this step is rarely taken. The objective of this kind of research is not so much music analysis as the development of music theory, or at least of computational music theory.

Similar arguments apply if the objective of analysis is to reveal the structure of cognition of a piece. Here the comparator which the computer output is to match is not a ready-made analysis, but structures in the minds of listeners, and direct comparison is not possible. Instead some kind of indirect test is applied: listeners are asked to respond in a specific way; researchers infer what response would be expected if the cognitive structure matched the computational one, and determine a degree of match between the two accordingly. (An example can be seen in Pearce and Wiggins, 2012.)

One needs to be aware that a perfect match is almost certainly unobtainable in either case, if for no other reason than that the human-made analyses often do not match each other and human perceptions differ. Again we are faced with a question of approximation: how close to human analyses and perception do computational analyses need to be in order to be confident that the computational procedure can be applied to as-yet unanalysed pieces and produce useful results? To answer this question we need some way of quantifying the divergence between human analyses or perception and computational analyses, and we need to know how good is good enough. Solutions to the first of these—means of measuring divergence—are commonly proposed, and form the basis of the annual MIREX competitions (Downie, 2008) (in which the topics are by no means all strictly music-analytical) where program outputs are compared with a "ground truth" (i.e., pre-existing data which is taken to be correct). The second desideratum, however—knowing how good is good enough—is rarely addressed and cannot be, at least not without some better understanding of how analyses are to be *used*.

1.5.2 Usefulness of Analysis

If we follow the reasoning of Nattiez outlined above (Sect. 1.3.2), questions of the 'correctness' of computational analysis seem irrelevant. The computational step produces only an 'inventory' and the real analytical step follows from the application of other factors to select from the inventory and interpret what it shows. In this case the question is not so much whether the inventory is correct, but whether it is sufficiently complete: does it contain the information required to make a correct analysis of the piece? Even if the answer is affirmative, we cannot be confident of a good analytical outcome because that does not depend on the completeness alone. One way to maximize the completeness of the inventory is to maximize its size, but this also increases the quantity of useless information and makes finding the useful information more problematic. More realistically, we must ask how much useful information is present in comparison to how much useless information. A perfect outcome would be for the inventory to contain only the information which the analyst selects to construct the analysis (in which case the analytical step of selection becomes trivial and redundant, and we effectively return to the situation discussed above where an analysis is judged ‘correct’ by the degree to which it matches an analyst’s conception). However, this perfection is unlikely to be achievable, and cannot reasonably be taken as a goal. Instead, if our objective is to use computer software as an aid in creating a ‘good’ analysis of a piece, and we will use our own judgement or that of some other expert to eventually determine a ‘good’ outcome, then our goal in software development should be a system which presents the user with an optimum which balances the risk of leaving out useful analytical information against the risk of obscuring the useful information by a heap of useless information.

If, on the other hand, we follow a conception of analysis as interpretation then we do *not* want the computational analysis to match a human analysis. Analyses are valued for their *difference* and for what is new that they can bring to our understanding of the piece analysed. It is highly unlikely, though, that any enriched understanding could come from merely random analyses. What bases are there for distinguishing a good analysis, meaning one capable of enriching understanding of a piece, from a poor analysis? (Indeed, how can we know whether or not someone’s understanding of a piece has been enriched?) Would the understanding-enriching principles be different from the music-theoretic principles used in finding ‘the structure’ of a piece, as envisaged in Sect. 1.4.3 above?

One possibility is that there are principles which we can apply to determine whether or not an analysis is useful, but there is no known procedure which is guaranteed to produce an analysis which is useful. Certainly this is the case if the principle is whether or not a human reader finds the analysis useful. We might perhaps be able to predict this with some degree of accuracy, but we can never be certain. This is not a reason to discount such a course of development though. On the contrary, I suspect that in the long run the manner of computational analysis which will prove most profitable for analysis (rather than for the development of music theory) is one which is interactive, presenting a human user with the results of computational analysis and allowing that user to modify or intervene in the procedure to arrive at an acceptable or interesting result. An example of such an approach is the ‘Automatic Timespan Tree Analyser’ which implements Lerdahl and Jackendoff’s method of analysis in interactive software (Hamanaka et al., 2006) (see also Chap. 9, this volume).

This and several earlier observations point to a potentially significant role in computational music analysis for machine learning, which could be applied in three ways. First, software might learn principles or parameter values from existing analyses (e.g., Pauwels and Martens, 2014). Second, following the idea that analysis seeks to uncover how a piece may be explained on principles which develop over the course of making the analysis (Sect. 1.3.1), the software might 'learn' the organization of the piece as it proceeds (see, e.g., Pearce and Wiggins, 2012). Third, the software might learn from the interactive input of the user. Fourth, the software might learn the probabilities of different musical configurations for use in an information-theoretic or minimum-description-length model of musical structure (see Chap. 7, this volume).

It was pointed out above (Sect. 1.4.6) that the authority of an analysis often derives from the expertise of the analyst. Machine learning has now advanced to the stage when computers can, within restricted domains, be *more* expert than humans. (In chess, for example, Garry Kasparov was beaten by IBM's Deep Blue computer in 1997, and more convincing computer victories against chess grand masters have followed.) This raises the intriguing prospect of computational music analysis which is, in some restricted sense, better than human analysis. Indeed, this appears to be the case already with software such as COSIATEC and SIATECCompress by David Meredith who presents evidence of the software finding patterns in pieces by Bach, Mozart and Chopin which can reasonably be claimed to be just as 'important' as other patterns in the same pieces identified by human analysts (Meredith, 2015).

Even in cases of super-human analysis, though, the extent to which the computer is doing 'music analysis' (as understood by musicologists) is uncertain. As suggested by the discussion above, we should probably re-imagine the objectives of the enterprise of computational music analysis. Instead of seeking a machine which analyses music automatically, we should think of computational music analysis as more like forensic science. Scientific examination of evidence can answer very specific questions, such as the likelihood that substances (e.g., DNA) found on the handle of a murder weapon came from the accused, with much greater accuracy than is possible without the application of scientific method. These answers can have a very significant role in court, but the scientific method does not ultimately determine the guilt or innocence of the accused; that continues to depend on the application of human judgement to the full range of evidence presented. Music-analytical computer software has advanced far beyond the stage of merely compiling an inventory of features or relations, as Nattiez envisaged, and can now answer complex questions about pieces of music. These answers are important and relevant for music analysis, but the final *musical* judgements, which determine how musicians might behave differently in future, will be made by people.

1.6 Conclusions

Music analysis is not a monolithic enterprise: different analysts do different things on different bases. Computational analysis therefore should also be multifarious. Most importantly, it should not and cannot be simply a machine reflection of the human activity. I conclude here with some reflections and recommendations about the ways in which computational analysis might profitably be used, and some recommendations on building software tools for music analysis.

1.6.1 Value of Computational Analysis

Computational music analysis needs to carve out a place for itself where it is not simply mimicry of human analysis, but a place which is not so distant from the human activity to prevent useful communication with musicians. We need to recall the potential value of computational analysis, the reasons we embark on this enterprise at all.

- Computational analyses have definite explicit bases, embodied in the analytical software used, whereas the bases for human analyses are, even when ostensibly explicit, subject to revision and reinterpretation.
- Computational analyses can handle large quantities of data, whether from within a single piece of music or from a large body of pieces of music.
- Computational analyses can formally test, try out or explore different hypotheses about musical structuring, without the risk of bias inherent in human analysis.
- Through techniques such as search and machine learning, computational analysis can find an ‘explanation’ for a piece which has a particular status among a well defined universe of alternative explanations, such as that it is the shortest possible, or most likely, description of the piece under certain assumptions.

The evidence that these can take music analysis into fruitful areas out of reach of human analysis has been emerging over the past decade or so. A clear sign is the two special issues of *Music Perception* dedicated to Corpus Methods (Temperley and VanHandel, 2013). In many of the papers in these two issues, conclusions were reported which would not have been possible without computational analysis. The level of sophistication in the analytical software used, however, was often at a much lower level than discussed in the contributions to this book. It would appear that we are still in the early stages of development of sophisticated analytical software capable of application to a large body of music to yield results in which we can have confidence. I suspect again that the reason is to do with error and approximation: research on a corpus usually involves conclusions based on statistics, and valid inference requires knowledge of the degree of error.

A second area in which computational analysis would show its distinctive value is by being embedded in other musical systems. As pointed out above (Sect. 1.4.5), music analysis rarely currently informs other musical activities. If an analysis really shows how a piece ‘works’, then could we not use that information to, for example, build a music recommendation system which offers us pieces of music which work well, or a system which adapts the music in a computer game to present coherently working segments of music aligned to the events in the game?

These possibilities might seem rather distant, but something similar is definitely within reach: the linking of analytical software to systems for rich visual display and sound output. Listening to music and reading an analysis are different activities, yet we do the second because we value the first. If we conceive of analysis as offering interesting interpretations of a piece of music (Sect. 1.4.5), then we should believe that analysis can influence our hearing of a piece, but for it to do so we have to somehow connect the analysis to the experience of listening. It is not at all self-evident how this should happen. To read something like ‘bars 72 to 96 are dominant preparation’ does not immediately translate to a way of hearing the piece. In illustration of this, I recount an event when I witnessed a visiting speaker at the university where I was a student. One of the staff of that university disagreed with the visiting speaker’s interpretation of the harmony of a passage, and in response the speaker did not use logical argument to persuade his adversary but instead played the passage again, playing some of the chords louder. To cause someone to hear a passage in a particular way, one needs to give them a listening experience. One of the potential values of computational analysis is that its outputs can be readily rendered in sound or in many kinds of visual analogy using the same computational resources as in the making of the analysis.

1.6.2 Analytical Tool-Building

Those who make analytical software come up against the ontological and epistemological issues of music analysis much more forcefully than do human analysts. The music expert can shift the analytical activity, select the material to consider and the conclusions to report, each in ways to soften the impact with ontological and epistemological problems. The computational researcher, by contrast, sets the software to work and cannot prevent it or its outputs from damage by such impact, except by once again rewriting the software, running it again and facing the same risks. It is for this reason that I believe computational researchers need to give these issues more explicit consideration than do human analysts.

- What is the input to the analytical software, and what is its status? Quantify, so far as possible, the differences between different possible inputs and the error inherent in them. What information comes from ‘the piece’ and what from elsewhere?
- What claim is being made about the analysis? Is it a structure with a particular status or property ‘within the piece’ (in which case be explicit about that status or property), or is it an image of the cognitive structures of the listener (in which case test it against psychological data), or is it some other kind of object?
- If the aim is to mimic some human analytical procedure, consider the degree of variance in human analysis. A perfect analysis from software of this kind is an impossibility if analysts do not agree in their analyses. Again, quantify error, and stop ‘improving’ the software when the error is close to the variance.
- If the aim is to offer possible interpretations of a piece, consider how the analyses will be presented to the user so that they can be usefully assessed. In particular, if the aim is to persuade the user of the benefit of hearing the piece in a particular way, use the computational resources to present the information in such a way as to make the transition from conceptual understanding to hearing more likely.

Finally, the objective of computational music analysis should probably not be to generate ‘an analysis’ but rather, like forensic science, to answer specific music-analytical questions with a degree of complexity, speed and accuracy which is impossible by other means.

References

Anglade, A., Benetos, E., Mauch, M., and Dixon, S. (2010). Improving music genre classification using automatically induced harmony rules. *Journal of New Music Research*, 39(4):349–361.

Bent, I. (1987). *Analysis*. Macmillan.

Byrd, D. (1984). *Music notation by computer*. PhD thesis, Indiana University.

Byrd, D. (2013). Gallery of interesting music notation. http://homes.soic.indiana.edu/donbyrd/InterestingMusicNotation.html.

Downie, J. S. (2008). The music information retrieval evaluation exchange (2005-2007): A window into music information retrieval research. *Acoustical Science and Technology*, 29(3):247–255.

Dunsby, J. (1982). Editorial. *Music Analysis*, 1(1):3–8.

Goehr, L. (2007). *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music*. Oxford University Press, second edition.

Good, M. (2001). MusicXML for notation and analysis. In Hewlett, W. B. and Selfridge-Field, E., editors, *The Virtual Score: Representation, Retrieval, Restoration*, volume 12 of *Computing in Musicology*, pages 113–124. MIT Press.

Goodman, N. (1976). *Languages of Art: An Approach to a Theory of Symbols*. Hackett, second edition.

Hamanaka, M., Hirata, K., and Tojo, S. (2006). Implementing "A Generative Theory of Tonal Music". *Journal of New Music Research*, 35(4):249–277.

Keller, H. (1965). The chamber music. In Robbins Landon, H. and Mitchell, D., editors, *The Mozart Companion*, pages 90–137. Faber.

Keller, H. (1985). Functional analysis of Mozart's G minor quintet. *Music Analysis*, 4(1/2):73–94.

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

Levinson, J. (1980). What a musical work is. *Journal of Philosophy*, 77(1):5–28.

Marsden, A. (2010). Schenkerian analysis by computer: A proof of concept. *Journal of New Music Research*, 39(3):269–289.

Mawer, D. (1999). Bridging the divide: embedding voice-leading analysis in string pedagogy and performance. *British Journal of Music Education*, 16(2):179–195.

McVicar, M., Ni, Y., Santos-Rodriguez, R., and De Bie, T. (2011). Using online chord databases to enhance chord recognition. *Journal of New Music Research*, 40(2):139–152.

Meredith, D. (2007). *Computing pitch names in tonal music: A comparative analysis of pitch spelling algorithms*. PhD thesis, Faculty of Music, University of Oxford.

Meredith, D. (2012). Music analysis and Kolmogorov complexity. In *Proceedings of the 19th Colloquio di Informatica Musicale (XIX CIM)*, Trieste, Italy.

Meredith, D. (2015). Music analysis and point-set compression. *Journal of New Music Research*, 44(3). In press.

Nattiez, J.-J. (1982). Varèse's 'Density 21.5': A study in semiological music analysis. *Music Analysis*, 1(3):243–340.

Nattiez, J.-J. (1990). *Music and Discourse: Towards a Semiology of Music*. Princeton University Press.

Pauwels, J. and Martens, J.-P. (2014). Combining musicological knowledge about chords and keys in a simultaneous chord and local key estimation system. *Journal of New Music Research*, 43(3):318–330.

Pearce, M. and Wiggins, G. (2012). Auditory expectation: The information dynamics of music perception and cognition. *Topics in Cognitive Science*, 4(4):625–652.

Pople, A., editor (1994). *Theory, Analysis and Meaning in Music*. Cambridge University Press.

Réti, R. (1962). *The Thematic Process in Music*. Macmillan.

Réti, R. (1967). *Thematic Patterns in the Sonatas of Beethoven*. Faber.

Samson, J. (1999). Analysis in context. In Cook, N. and Everist, M., editors, *Rethinking Music*, pages 35–54. Oxford University Press.

Sapp, C. (2014). 371 Four-part Chorales by J.S. Bach in the Humdrum file format. https://github.com/craigsapp/bach-371-chorales.

Sleator, D. and Temperley, D. (2003). The Melisma music analyzer. http://www.link.cs.cmu.edu/melisma/.

Sturm, B. L. (2014). The state of the art ten years after a state of the art: Future research in music information retrieval. *Journal of New Music Research*, 43(2):147–172.

Temperley, D. and VanHandel, L. (2013). Introduction to the special issue on corpus methods. *Music Perception*, 31(1):1–3.

van Kranenburg, P. (2008). On measuring musical style—The case of some disputed organ fugues in the J. S. Bach (BWV) catalogue. *Computing in Musicology*, 15:120–137.

Whorley, R. P., Wiggins, G. A., Rhodes, C., and Pearce, M. T. (2013). Multiple viewpoint systems: Time complexity and the construction of domains for complex music viewpoints in the harmonization problem. *Journal of New Music Research*, 42(3):237–266.