In: Music in the Social and Behavioral Sciences, An Encyclopedia Edited by:B. Thompson. SAGE (forthcoming)

Computer Assisted Music Analysis

Christina Anagnostopoulou[1]
Faculty of Music Studies, University of Athens
chrisa@music.uoa.gr

The use of computers, and of the science of informatics in general, has penetrated into many fields of study related to music. One of the most important fields amongst these is the study of the musical object itself, i.e. music analysis. Computer assisted music analysis is becoming more and more widespread, as computers can cope with extensive amounts of data, do comparative analyses of large musical corpora, and formalise the analytical process so that there are no gaps or hidden steps in the procedure. More specifically, computers are particularly good at categorising musical works, or musical segments, and at finding repeated patterns that occur in one or more works. Computer assisted music analysis forms part of the larger field of computational musicology, and has close links with the area of music information retrieval (MIR). Perhaps the most interesting challenge in this enterprise is to combine the objectivity and scientific rigour of informatics with the interpretative nature of music analysis.

Aims of computer assisted music analysis

There is an obvious answer to the question of why use a computer to analyse music: A computer can deal with large amounts of data, and make calculations that a human might not be able to do, or that might be tedious, or take a very long time. Although this might be a valid reason, it is by far not the only one: More importantly, a computer nowadays can give answers that a human might not have been able to find at all, as it can learn from data and draw novel conclusions. For example, the computer is very good at making comparisons, and therefore many times it is used to create comparative analyses of large musical corpora (such as, for example, Bach chorales with chorales of other composers, or folk songs of different countries) and spot large-scale patterns that would not have been possible to detect otherwise.

The main objective in computer assisted music analysis is to produce musicologically interesting results. This might sound obvious, however, it might not always be straightforward: Music analysis should not be a mere 'output' of an algorithm; the role of the analyst in this process is indispensable in evaluating the results and rethinking the algorithmic procedure. Another aim, just as important, is the formalization of an analytical process, since music analysis is a human task which could be formally modeled, making sure that there are no hidden steps and underlying assumptions in the procedure. The analyst always needs to take certain decisions, especially regarding analytical criteria, which in music are most times context dependent. There are many degrees of automation and assistance, and in this context, a computer-assisted analysis is just as significant, and perhaps even more meaningful, than a fully-automated one. Often, in practice there is another aim which comes from Computer Science: to test computational methodologies in a very challenging and abstract domain, such as music. Finally, computational analysis can be cognitively inspired, or model specific cognitive precesses related to music, such as for example David Temperley's work (see reading below).

---

1 The paper appears in C Anagnostopoulou (2014) Computer Assisted Music Analysis In: Music in the Social and Behavioral Sciences, An Encyclopedia Edited by:B. Thompson. SAGE

The musical input and music representation

The input to the computer in order to perform a music analysis task can be either the raw audio signal, or some encoding of the score which employs discrete pitch and rhythm information. Analysis from the audio usually relates to low-level features of the sound, such as the signal's energy and zero-crossing rate, various spectral descriptors, the fundamental frequency, the chroma vector, and others. These features are commonly used in pattern recognition and discovery applications which are capable of generating higher-level signal representations, including tonal, harmonic and structural descriptions.

The score, in order to become an appropriate input to computational analysis, can take various formats, such as MIDI (www.midi.org), MusicXML (http://www.musicxml.com/), Kern (http://www.music-cog.ohio-state.edu/Humdrum/representations/kern.html) and others. In the analysis from the score, since the information of discrete pitches already exists, more high-level musicological tasks can be performed. For example, more advanced harmonic and structural analysis, motivic and semiotic analysis, reductionist approaches to analysis, tonal analysis, and others.

Whether sound or score, the input has to be translated into musical knowledge before any computation takes place. This is the issue of music knowledge representation, a field in itself, which looks at what musical knowledge is extracted and how it is represented in order to be useful to further analysis. Music is structured at multiple levels, such as rhythmic, melodic, and harmonic, and each of these dimensions can be described in many different ways. For example, melodic intervals and melodic contour might need to be explicitly represented. There are various representational schemes; One of the more general, clear and frequently used ones is the Multiple Viewpoint Formalism, initially developed in the mid-nineties and advanced several times till today (see further reading below). In this representational scheme a crucial component is that each musical parameter can be independently represented. The intuition behind multiple viewpoints is that no single music representation can be sufficient for music, and that using a number of different viewpoints can give more interesting analytical results.

Types of tasks performed by a computer

Repetition, variation and transformation are inherent musical processes, and most music analysis methodologies use comparative techniques in order to study relations of similarity and contrast between musical components of a musical work or a number of works. At the same time, computers can be particularly good at two related tasks: Categorisation based on similarity, and repeated pattern discovery. Categorisation consists in classifying musical works, musical sections or musical segments of varied length into categories according to how similar they are. For example, attributing musical works to composers or musical genres, or categorising musical segments to paradigmatic classes (as in the case of Paradigmatic Analysis).

Pattern discovery requires initially the explanation of a pattern in music. Patterns in music can be any sequences of notes (i.e. musical segments), or sequences of values of musical parameters, such as for example sequences of rhythmic values taken from a score, sequences of intervals, sequences of melodic contour directions, and so on. Therefore the concept of pattern is particularly close to that of a motif, segment or phrase in music, whether this consists of notes, or of rhythmic values, or of intervals, or other parametric values. The discovery of such patterns, which repeat either as identical or varied, especially in large musical corpora is a type of motivic analysis which is often necessary in music. Related to the above are other analytical tasks which can be carried out computationally, such as segmentation, harmonic analysis, semiotic analyses, reductionist analyses, tonal analyses, and so on.

The machine learning paradigm and music information retrieval (MIR)

In the last decades, a new way of thinking has been put forward in computer science, and artificial intelligence in particular: machine learning. Machine learning algorithms are able to learn from data, rather than be told explicitly what they need to know or what to do. Often training is needed in order to learn, just like it is the case with humans. These algorithms can perform various tasks such as classification of pieces into genres, or smaller structures into any type of classes, which is a central task of music analysis. Evaluation of classification methods is usually according to the accuracy of class label assignment of new pieces or musical segments / structures. Machine learning algorithms have thus been very popular in recent years in the domain of music analysis, and a new field of study has emerged, studying large musical corpora for mainly commercial applications, which often uses machine learning algorithms: Music Information Retrieval (MIR). Conferences such as ISMIR (International Symposium of Music Information Retrieval, see http://www.ismir.net/) include many such cases. Music information retrieval is possibly the largest area of research in computer approaches to music today. The objectives of this research are not the same as in music analysis. However, the boundaries between the two are often not clear, and MIR produces research which is related to music analysis, especially in providing methods and results which can often be further analysed to produce an interesting computational music analysis study.

Examples of music analysis by a computer

There have been several systems which perform computer assisted music analysis. One of the most well-known ones is HUMDRUM (http://www.musiccog.ohio-state.edu/Humdrum/), which takes as input the Kern representation format, a text file which represents the musical score, and produces basic calculations on the score, such as interval and contour calculations, as well as some more advanced structural interpretations.

Anthony Pople's Tonalities project and related software has been another historically significant example. Pople presented a broad theory to describe tonality, and created a system which could analyse up to early atonal music, using various concepts related to tonality, including complex set theory.

Other examples include Olivier Lartillot's approach to motivic analysis (Lartillot, 2008), Alan Marsden's system for Schenkerian analysis (Marsden, 2010), the approaches described in the dedicated Special Issue of the Journal of Mathematics and Music, 4(2), and in the related issue of Computing in Musicology, 14, entitled Music Analysis East and West (see suggested reading below). Examples of music classification can be found in abundance in each year's ISMIR proceedings (www.ismir.net).

The formal analysis of folk music, which as a starting point one could consider Alan Lomax' Cantometrics, is gaining ground in recent years, as folk music corpora seem ideal for computational music analysis. This work has resulted in an area of study of its own, known nowadays as Computational Ethnomusicology.

Explaining and evaluating results

The results produced by any analytical algorithm need to be interpreted and evaluated by a human analyst in order to be musically meaningful. Often the process is circular, which means that the observations and evaluations of the analyst go back to fine-tune the representation, segmentation (where it exists), and algorithmic procedure, in order to achieve more informed results. The use of statistics in the last decades can produce results which are statistically significant, but again their musical validity needs to be addressed. As analysis is a mainly interpretative task, the computer can suggest solutions, but the role of the analyst in their evaluation is vital in order for analysis to keep its interpretative nature. Related to this, a general question which is implicit in the field of computational music analysis is whether we will ever be able to create intelligent systems that can analyse music without any intervention from the analyst. The question is still open, and a challenge for future research.

Christina Anagnostopoulou
University of Athens

See also: Computer models of music, systematic musicology, music analysis, computer music, algorithm.

Further readings:

Marsden, Alan. "What was the question?" Music Analysis and the Computer. In *Modern Methods for Musicology*, Ed. T. Crawford and L. Gibson. Ashgate, 2009.

Anagnostopoulou, Christina and Buteau. Chantal. Can computational music analysis be both musical and computational? *Journal of Mathematics and Music* v.4/2 (2010). This whole special issue contains examples of computer assisted music analysis.

KERN Representation: http://www.music-cog.ohio-state.edu/Humdrum/representations/kern.html (accessed September 2013)

Huron, David, "The Humdrum Toolkit: Software for Music Research" (n.d.) Available at http://www.musiccog.ohio-state.edu/Humdrum/ (accessed September 2013)

Pople, Anthony. Modeling musical structure. In: *Empirical Musicology – Aims, Methods, Prospects*, OUP 2004.

Cook, Nicholas. Computational and Comparative Musicology. In: *Empirical Musicology – Aims, Methods, Prospects*, OUP 2004.

Hewlett, W.B. and Selfridge Field, Eleanor. *Music Analysis East and West: Computing in Musicology 14,* MIT Press, 2006.

Conklin, Darrell and Witten, Ian. Multiple Viewpoint Systems for Music Prediction. *Journal of New Music Research* 24(1), (1995). The work is continued at: D. Conklin, Multiple Viewpoint Systems for Music Classification. *Journal of New Music Research* v.42/1 (2013).

Lartillot, Olivier. *Motivic pattern extraction in the symbolic domain*, in Intelligent Music Information Systems:Tools and Methodologies, J. Shen, V. Shepherd, B. Cui, L. Liu, eds., IGI Global, 2008.

Tzanetakis, George, Kapur, Ajay, Schloss, W. Andrew and Wright, Matthew. Computational Ethnomusicology. *Journal of Interdisciplinary Music Studies*, v.2/1 (2007).

Marsden, Alan. Schenkerian Analysis by Computer: A proof of concept. *Journal of New Music Research* v.39/3 (2010).

Cambouropoulos, Emilios. The Musical Surface: Challenging Basic Assumptions. *Musicae*

Scientiae, Special Issue (2010)

Temperley, David. *The Cognition of Basic Musical Structures*. MIT Press, 2001.