Music, Intelligence and Artificiality

Alan Marsden

in E.R. Miranda (ed.), *Readings in Music and Artificial Intelligence* (Amsterdam: Harwood, 2000)

Abstract

The discipline of Music-AI is defined as that activity which seeks to program computers to perform musical tasks in an intelligent, which possibly means human-like way. A brief historical survey of different approaches within the discipline is presented. Two particular issues arise: the explicit representation of knowledge; and symbolic and subsymbolic representation and processing. When attempting to give a precise definition of Music-AI, it is argued that all musical processes must make some reference to human behaviour, and so Music-AI is a central rather than a peripheral discipline for musical computing. However, it turns out that the goals of Music-AI as first expressed, the mimicking of human behaviour, are impossible to achieve in full, and that it is impossible, in principle, for computers to pass a musical version of the Turing test. In practice, however, computers are used for their non-human-like behaviour just as much as their human-like behaviour, so the real goal of Music-AI must be reformulated. Furthermore, it is argued that the non-holistic analysis of human behaviour which this reformulation entails is actually informative for our understanding of human behaviour. Music-AI could also be fruitfully concerned with developing musical intelligences which were explicitly not human. Music-AI is then seen to be as much a creative enterprise as a scientific one.

Introduction

Computers are machines. Intelligence is a human characteristic, and though it is often taken to be the characteristic which distinguishes us from animals, computers rarely approach the intelligence even of animals. One of the characteristics of machines is that they are man-made, and in that sense artificial (the sense of artificial as “unreal” will be discussed briefly below). The characteristic which distinguishes them from other artificial things is *behaviour*, and this characteristic is one which they share with humans and animals. In fact everything has behaviour, in that everything responds in a particular way in interaction with an environment: if one drops a ball it bounces; if one drops a glass it smashes. The real distinction between machines and other artificial objects cannot be made without reference to human values and “intentions”: we value machines because of their behaviour and not because of other characteristics (e.g. their shape, dimensions and solidity, as in the case of chairs). We use machines to extend our own behaviour. A class of machines which has become particularly important during this century is machines whose behaviour concerns information. This class contains such ancient machines as the printing press and such common ones as the telephone — it is a mistake to regard “information technology”, at least in this sense, as something new.

The characteristic which computers have which is genuinely new, and which sets them apart from other information-processing machines, is that their behaviour is not only controllable by the user (this is an important characteristic of all useful machines) but that their behaviour is *definable* by the user. Other machines can have this characteristic, both information-processing machines and others, but only within tight constraints. In the case of a computer, on the contrary, its behaviour is highly unconstrained, at least in the domain of the processing of information (the possibilities for physical behaviour are usually very limited). In fact, the ideal computer is a "universal processing machine" which is capable of performing *any* kind of behaviour in the domain of abstract information processing. At the level of programming, the "input" which a computer reads is a definition of a kind of behaviour, or in other words, a definition of an abstract machine. If computers are thus intended to be able to mimic any kind of behaviour, it is not surprising that there should be interest in programming computers to behave in ways that are human-like and which could be called intelligent. There has also been interest in having computers perform musical tasks, whether it be playing music, processing music, or creating music. Whether behaving in a musical manner implies behaving in a human manner is discussed below. For now, it is sufficient to note that the combination of the two — the intention to behave in a human-like fashion and to perform a musical task — is the topic of this paper.

History

An argument is presented below that no attempt to have a computer perform a musical task can be totally unconcerned with the issues of Artificial Intelligence, but customarily Music-AI has included only those musical computer systems which have involved a degree of complexity which is not the complexity of mathematical formulae, nor the complexity of large quantities of data, but rather a kind of complexity of ideas. As in other domains, certain tasks have been considered to involve intelligence while others have not. (This is a problematic issue, which will be returned to below.) Sound synthesis, for example, is an area which has attracted a great deal of very successful work, but little of it is regarded as being in the domain of Music-AI because it has concerned acoustic and psycho-acoustic phenomena and the mathematics of signal processing rather than being concerned with thinking. Similarly, the vast area of systems for capturing, processing and using performance data via sequencers and the like is excluded from the domain of Music-AI, as are systems for music notation. A brief historical survey is presented here, organised around different architectures of system.

Early attempts at programming computers to perform musical tasks took an algorithmic approach. The objective was to describe the procedures which must be performed in order to produce a musical result. An example of high-quality work of this kind can be found in the research of Longuet-Higgins and his co-workers (Longuet-Higgins, 1978; Longuet-Higgins & Steedman, 1971). The objective of this work was a system which could transcribe music played on a keyboard (the work began in the days before MIDI) to music notation. This involves resolving issues about the representation of pitch (should a note be written as C sharp or D flat, for example), which involves determining key, and issues about the representation of rhythm, which involves both determining metre and coping with the variations from metronomic playing of a real performance (which can be quite severe). Algorithms with a moderate to high degree of success for these tasks were designed and implemented in the language Pop-11. Projects which have also taken an algorithmic approach have been directed at tasks as diverse as composition (e.g. Ames & Domino, 1992; Cope, 1991) and transcription of lute tablatures (Charnassé and Stepien, 1992).

While this approach can produce good results, if those results are to be applicable in other programs to perform other tasks, then it is up to the researcher to make certain that the algorithms are suitably designed and explained. Some authors, (Longuet-Higgins among them) are excellent at explaining what their algorithm does; others are not so. The algorithms themselves, without explanation, cannot be expected to be transferable to a program to perform another task, however similar. At issue here is really the nature of the principal objective of research in Music-AI. Is it to design and implement computer systems which perform musical tasks (an engineering objective), or is it to discover and explain the knowledge which underlies these tasks (a cognitive-science objective)? Most researchers would claim the latter, but this can only be tested by achieving the first objective to some degree also.

While every computer program ultimately comes down to algorithms, there has been considerable interest in devolving the translation from knowledge to algorithm to the computer so that the representation in which a system is expressed can be more directly a representation of the knowledge underlying a particular task. A number of formalisms intended to achieve this have been designed. The one most often been used in music, usually because of a perceived similarity with language, has been formal grammars. Another early example of Music-AI is the harmonic analysis system of Winograd (1968). The core of this was a systemic grammar which described the configurations of chords, harmonies and tonalities possible in homophonic tonal music such as the chorale harmonisations of J.S. Bach. This gave an extremely clear exposition of the “knowledge” of tonal theory. The grammar could be applied in the analysis of a piece of music to discover how the grammar accounts for the piece, and thereby, by reporting the steps of the derivation, producing a harmonic analysis of the piece. However, many different analyses were possible for any one piece (musicians will be familiar with the idea of different possible analyses, but they might be surprised at quite how many were allowed by Winograd’s grammar, which was quite a faithful reproduction of classical tonal theory.) The part of the system which derived analyses, therefore, called the “parser”, had to be quite complex and make use of other, procedural, “knowledge” in order to arrive at harmonisations which were acceptable. In principle a grammar should be applicable in either direction, i.e. to either analyse music of produce music. It might be possible to use Winograd’s grammar to produce harmonisations, but Winograd did not attempt this. A well-known grammar which did produce music was that of Baroni et al. (1992), who produced a number of grammars to generate chorale melodies, eighteenth-century French chansons, and the text repetition patterns of Legrenzi arias. In both Baroni et al. and Winograd’s work, the business of translating the grammar to an algorithm was not devolved to the computer, as suggested above, but coded by hand. In the case of Kippen & Bel’s Bol processor (1989, 1992), however, the computer system operated directly on the grammars. The Bol processor was a system intended to assist in the understanding of a style of *tabla* drumming found in North India. It was capable both of producing new pieces of music, and of analysing existing pieces. Their publications also include excellent discussions of the principles of using grammars in this kind of work and of some of the issues involved.

Another formal systems for representing knowledge applied in Music-AI is KL-ONE, a well-developed system of knowledge representation, derived from frames and semantic nets, which expresses knowledge in terms of concepts and roles, and defines inheritance and other relations between them. Here again, the intention is to allow a clear expression of knowledge which is susceptible to direct implementation by computer. Furthermore, this knowledge is, in principle at least, expressed abstractly without any reference to its application in any particular task. KL-ONE is used to provide the symbolic layer of HARP, a hybrid system applied to a number of musical tasks, often involving real-time interaction between a performer and a music-production system (Camurri et al., 1994).

One of the problems Kippen & Bel identified in developing their Bol-Processor grammars was the difficulty of knowing what should go into a grammar: how is the researcher to determine what the rules of the grammar should to be? The common paradigm has been to make a first attempt, to examine its results, then, on a rather *ad hoc* basis, to attempt some revisions to the grammar which will correct the errors of the previous results. The cycle of testing and revision then begins again. Such a strategy will probably never produce a perfect system, though it might approach perfection, but the *ad hoc* nature of the rule revision is disconcerting: how can the researcher have any confidence that the revisions are the best to propose in the circumstances? It is a characteristic of an intelligent animal that it *learns* from its experience and performs better next time in similar circumstances. In fact, this behaviour is more characteristic of intelligence than is behaving well in every circumstance. One of the goals of Artificial Intelligence, then, is systems that learn, and these can be found in Music-AI also. Kippen & Bel attempted to build learning into their system so that rule strengths could be adapted automatically and so that at least some of the new-rule generation process could be automated (1989). Musical learning systems, however, are best exemplified in the work of Widmer, who has completed projects which learn counterpoint rules (1992) and which learn expressive performance (1996). Cope's EMI system (1991), which learns to compose music in the style of the music given to it, does not properly belong in this category of intelligent learning systems because the learning requires a considerable degree of input from the user of the system also. While it is true that intelligent animals often learn best with teachers, these teachers do not interfere with the functioning of the animal in any way other than the normal channels of interaction. (Teachers do not resort to brain surgery, in other words.) Furthermore, it is a characteristic of intelligent animals that they learn *spontaneously*, and it is this characteristic that is most sought in AI research in learning.

A number of characteristics of intelligent behaviour, including the one of spontaneous learning just mentioned, gave rise at the end of the 1980s to a new paradigm in computing variously called connectionism, parallel distributed processing, and neural networks. Two of the most important motivations were the observation that intelligent behaviour could not possibly arise from the mechanisms proposed by traditional "sequential" AI approaches at the speed at which it does in animals. Furthermore, it is a characteristic of intelligent animals that, in surroundings which they have never before encountered, and therefore surroundings for which they have no perfectly applicable knowledge, they are able to perform tolerably well. Traditional AI systems, however, when presented with something somewhat different from their intended task, generally perform spectacularly badly. This is sometimes referred to as "brittleness". In the new paradigm, which is clearly explained in Leman (1992) and other sources, the behaviour of a system results from the net effect of the behaviour of a number, possibly a very large number, of simple but interacting processing units. When appropriately configured, such systems are capable of learning, in the sense that their behaviour approaches the desired behaviour. Furthermore they typically perform moderately well with unfamiliar input rather than exhibiting the brittleness of classical systems. Such systems have been used with remarkable success in such diverse domains as tonal theory (Leman, 1994, 1995), the classification of timbre (Cosi et al., 1994), and the quantisation of rhythm (Desain & Honing, 1992). Desain & Honing (1992) include a direct comparison of a classical and a network system performing the same task. From the engineering perspective, such systems often perform well. From the cognitive-science perspective, however, they involve a total shift of philosophy. It in inappropriate to use a network system in the hope of discovering the rules of tonal harmony, for example, at least in the form that they are traditionally expressed. The "knowledge" which a network system acquires during its learning is distributed through the connections of the network; one cannot necessarily examine the state of the network after training and directly extract from it a rule in the form "if X then Y", as one often can from a learning system based on classical computing.

The philosophical shift has justifications other than the utility of such network systems, expressed in Leman (1993), Lischka (1991) and Kaipainen (1996), but it is important to realise quite how different it is from the cognitive science which gave rise to grammars, KL-ONE, and the like. Nor should it be thought that the new paradigm has supplanted or should supplant the former one. Much recent work involves both kinds of computing (e.g. Camurri & Leman, 1992 and Goldman et al., 1995), often assigning "subsymbolic" processing to a network while "symbolic" processing is carried out using a more traditional kind of architecture. However, care must be taken in ensuring that the mixture of the two philosophies is sound if the goal of improving understanding musical behaviour — the cognitive-science goal which was argued above to the fundamental to Music-AI — is not to be compromised.

Philosophy

In a precise discussion of Music-AI, there are three terms to be defined: "music", "artificial" and "intelligence". Some adumbrated definitions were given above. "Artificial", for example, was taken to mean man-made and not occurring naturally in the universe. By this definition music is also artificial, as is any other human product. A tightening of the definition is warranted, restricting the word "artificial" to refer to human products which are intended to emulate something else (which probably, but perhaps not necessarily, occurs naturally), hence artificial pearls, etc.

"Music" is notoriously difficult to define (for a straightforward discussion of some of the issues, see Davies, 1978), but all agree that while it involves sound, it is impossible to define solely in terms of sound. The classic test case is John Cage's *4'33"*, during the performance of which the only sounds heard are those which happen to occur in the environment — the performer is not instructed to make any sounds at all. If this piece, in which any sound can occur, is to be taken as music, then any sound is music and so all sounds are music. This is clearly unsatisfactory as a definition of the word as normally understood. Even if this extreme case is not admitted as a piece of music, it is not difficult to name pieces in which all kinds of normally non-musical sounds have been included, and it is extremely difficult to find physical differences between the sounds which characterise music and those that do not. Thus definitions of music generally make reference in some way or other to human activities, whether composition, performing or listening. If, then, the very definition of music requires reference to human activities, any computing system which is supposed to perform a musical task must also take account of those human activities. As an example, consider a sound-synthesis system, a common kind of musical computing system which is not normally considered an example of artificial intelligence. In designing any such system, choices must be made about the frequency ranges to be accommodated (and hence the sampling rates to be used). For a *musical* system, the appropriate choices are to set the frequency range to the maximum humanly audible range, since the results are intended to be listened to by people and not bats or any other animal with a different audible range. Pursuing the example further, suppose that the designer wishes the user of the system to be able to specify the sound output in terms of individual sound events, which we might call “notes”, and to specify the time of occurrence for each note. This will require some reference to the phenomena by which we segment a stream of sound into separate events, and also an understanding of where the perceived “start-time” of a note is in relation to the physical beginning of the sound, its amplitude envelope, etc. Going yet further, the user might want to be able to specify the grouping of notes into phrases and have this phrasing reflected in the synthesised sound. This would require an understanding of the relation of variations in timing and other factors to perceptions of phrase beginnings and endings (see Todd, 1985; Sundberg, Friberg, & Frydén, 1991). The point of the argument is that if any system is to be *musical* it must make reference to human behaviour, and to that extent *any* musical system must involve artificial intelligence. There is no obvious place at which to draw a boundary between where one must take into account human behaviour which is not intelligent, and where one must take into account behaviour which is intelligent. By this argument, furthermore, the discipline of Artificial Intelligence becomes not a peripheral specialisation but a core element of successful computer science.

“Intelligence” is the most difficult of the three terms to define, and the one whose definition is most contentious. It was suggested in the introduction above that artificial intelligence meant programming computers to behave like people. Later, spontaneous learning and performing with moderate success in unfamiliar surroundings were suggested as characteristics of intelligent behaviour. A third definition is suggested by a common usage of the word “intelligent” with respect to software. An “intelligent help system”, for example, is one which determines the information to be provided to the user on the basis of the user’s recent activities. In other words the behaviour of the system is sensitive to its environment. This is true of every piece of software — its output it determined by its input — but here there is a significant difference in the domain of the input. Normally software uses a very restricted input; so-called “intelligent” software instead attempts to receive input from as much as possible of its environment. Clearly this is related to the definition of intelligent behaviour as performing moderately well in unfamiliar surroundings, because attention is payed to the totality of the surroundings. Furthermore, if the environment is taken to include the past, then this definition of intelligence as behaving appropriately in the environment will include learning also. However, computers generally have extremely limited channels for receiving input from the environment, and considerable work is needed in this area if we are to see behaviour which is really intelligent under this definition. In fact, if we really want an intelligent computer to behave in the same way in which a human would in a given environment, including that environment’s past, then the computer would have to have the same channels of input, the same memories, the same means of acting upon the environment, and indeed the same objectives. In short, the computer would *be* that person. Artificial intelligence under this definition, then, is an impossible goal.

Some of these difficulties are overcome by limiting the channels of communication, as in the definition of intelligence encapsulated in the “Turing test”, proposed by Alan Turing at the very beginning of the discipline of Artificial Intelligence. The test is as follows. Two rooms have teletypes (the technicalities are not significant — any restricted means of communication usable by both computers and humans would do) as the only means of communication with the outside world. In one room is a computer connected to the teletype, in the other a person. Those on the outside may ask questions via the teletype, in a restricted domain. If they cannot tell from the responses to the questions which room contains the computer and which the person, the computer has passed the test and may be described as intelligent. A musical version of this test could be proposed also. (For a similar argument making a point related to the one above about the essentially human nature of musical activity, see Cross, 1993.) Two rooms are set up with a channel by which music is communicated to the outside world. We might also allow a channel by which some sort of feedback (applause, perhaps, or other pieces of music) goes into the room. In one room is a composer; in the other is a computer. The test is passed when those outside the rooms cannot tell which contains the computer. While it might be possible for a computer to pass this test in practice (i.e. in an empirical sense), there is an argument that a computer could never pass the test in principle (i.e. in a rationalist sense). (While the test might appear inherently empirical, because it fundamentally involves observations, it is generally not conducted in practice but as a “thought experiment”, and so is not empirical at all.) It is often argued that originality is an essential characteristic of music. (From the perspective of composing, this is commonplace; for a perspective from listening, see Kunst, 1978). Computers are digital automata, and so their behaviour is always, in principle at least, predictable and therefore *cannot* be original. Thus a computer cannot, in principle, pass this test. There is a persuasive counter-argument that dynamic systems, and so-called chaotic systems in particular, can be deterministic, in the sense that their future state is entirely determined by their current state, but yet unpredictable. In fact such systems have been used for creating both music and visual art (the visual examples are quite well known; see Little, 1993 for a musical example). However, this depends, in principle, on the dynamic system operating in an infinite domain (e.g. using rational numbers), and computers can only simulate this by a finite domain of very many elements. The argument in principle, therefore, remains. The argument in practice will not be defended because clearly it is a hopeless task for a person to know all the details of the state of a computing system, finite as the number of possible states might be. Indeed, it is now to matters of practice that we will turn.

Pragmatics

If the goals of Music-AI suggested above — behaving in a completely human-like way and composing music indistinguishable from humanly-composed music — are impossible to achieve, what should Music-AI realistically aim at? In fact, we frequently want superhuman, and therefore non-human, behaviour from computers. We often want computers to process data in larger quantities, at greater speed and with greater accuracy than is humanly possible. In these cases, putting aside questions about whether the computer's behaviour is really intelligent, it is precisely because it is artificial (other-than-human) that it is useful. Thus the real goal in developing a computer system is often for it to behave in a human-like manner in some respects but in a non-human-like manner in other respects. There are two difficulties of definition here. Firstly, the boundaries must be defined: in which respects is the behaviour to be human and in which respects is it to be super-human? Secondly, there is a difficulty in knowing what human behaviour is when constrained in the appropriate respects. We can realistically know what human behaviour is in total by observation. To observe only certain respects we are in danger of either getting a false picture or of distorting the behaviour so that it is no longer truly human. To take a concrete example, consider the case of designing a system which is to transcribe musical rhythms into notation. (For actual systems of this kind, see Lee, 1985 and Desain & Honing, 1992.) What is the human behaviour which the system is to mimic? We could attempt to answer this question by examining musicians' transcriptions of actual music. But real music contains details other than rhythm, and these almost certainly have an influence on the transcription. We could expand the task domain of the system so that some of these other factors are taken into account, but, by the argument above, there will always be factors not taken into account. On the other hand we could attempt to answer the question by having musicians transcribe material in which these other details were neutralised (e.g., all pitches could be the same). Now, however, the musicians would not be dealing with real music and so would be performing an unreal task. Thus we are caught between having to deal with only partial information or information which is of doubtful validity. The behaviour of the computer system is thus artificial also in the sense that even the supposedly human-like part of its behaviour is not really human behaviour.

Although this has been described as a problem (and it is a problem also for psychology), it can be turned into a virtue. While the arguments above would suggest a holistic approach to the study of human behaviour, it can be argued that the understanding we seek is not holistic but analytical. Indeed if we are to design computer systems which behave sometimes like humans and sometimes not, we need to have an analytical, piece-by-piece, understanding of human behaviour. While the artificial division of human behaviour described above is anathema to a holistic approach, it is an essential tool in an analytical approach. Anyone taking this approach has to accept the artificiality, but its effects can be reduced, for example, by taking "slices" of human behaviour in different ways and attempting to unify the results, thus neutralising the effects of ignoring relevant factors or of studying an artificial task. (The rhythm-transcription study alluded to above, for example, might examine the results of *both* of the kinds of investigation suggested there, and it might also take into account data from studies of the grouping of notes in the pitch domain, since grouping is probably a factor in rhythm transcription and we might presume that common grouping behaviours are used in both domains.) But, on the other hand, the artificial division of human behaviour can be regarded as a tool to reveal the component details of behaviour, like a kind of dissection. It is only when we impose an artificial task that we can begin to break down behaviour into manageable pieces. The arguments here are closely related to those between network computing systems, which operate in a holistic manner and suggest a holistic approach to behaviour, and symbolic systems which operate in a step-by-step manner and suggest an analytical approach.

Finally we must consider the possibility that artificial intelligence is not necessarily a copy of human intelligence. The definition of intelligence as “behaving appropriately in the current environment”, discussed briefly above, does not make any necessary reference to humans, so long as we can define “appropriately” without reference to humans. Appropriate behaviour might be defined, for example, as behaving in a manner which leads to a certain goal. Again, most software does this; the difference in “intelligent” software is that its response to the environment is more perspicuous and its means of achieving goals more adaptive. This suggests the possibility of computer systems which behave intelligently, but which do so in a manner quite different from our own: artificial intelligence in which “artificial” means unreal and unusual as well as man-made. For musicians this is an attractive possibility, since it suggests the generation of entirely novel approaches to music, but approaches which, because they are intelligent, are interesting and productive. The most obvious application of this kind of system would be in composition, whereby entirely novel kinds of composition could be generated. (See Laske, 1989, for further discussion.) There could equally well be applications in analysis, where an objective is often to arrive at a novel understanding of a piece of music, whereby unforeseen approaches could arise. The same applies also to performance.

By both of the last two arguments, the argument that artificial intelligence is an analytical tool for the study of behaviour, and the argument that artificial intelligence can be explicitly non-human, Music-AI becomes a kind of creative enterprise rather than a purely scientific one. Designing explicitly non-humanly intelligent systems is obviously a creative task. Using artificial intelligence as an analytical tool, however, is less clearly so. We only need recognise, though, that at each stage of the process of developing a system we are making choices: when deciding on the domain of the task; when deciding how to circumscribe and observe the human behaviour to be mimicked; when combining data from different kinds of observations. Overall, one is *creating* a perspective on, and model of, human behaviour. Neither is it immediately obvious that the criteria by which the perspective and model are to be judged are scientific ones and not artistic ones. It is not so bizarre, then, to couple the artistic realm of music and the science of computation into the enterprise of Music-AI.

References

Ames, C. & Domino, M. Cybernetic Composer: an overview. In M. Balaban, K. Ebcioglu, & O. Laske (eds.) *Understanding Music with AI: Perspectives on Music Cognition* AAAI Press and MIT Press (1992) 186–205.

Baroni, M., Dalmonte, R. & Jacoboni, C. Theory and analysis of European melody. In A. Marsden & A. Pople (eds.) *Computer Representations and Models in Music*, Academic Press (1992) 187–205.

Camurri, A. & Leman, M. Hybrid representation of music knowledge - a case study on the automatic recognition of tone centers. Proceedings of the International Workshop on Models and Representations of Musical Signals, Capri (5–7 October 1992).

Camurri, A., Frixione, M. & Innocenti, C. A cognitive model and a knowledge representation system for music and multimedia. *Journal of New Music Research*, v.23 (1994) 317–347.

Charnassé, H. & Stepien, B. Automatic transcription of German lute tablatures: an artificial intelligence application. In A. Marsden & A. Pople (eds.) *Computer Representations and Models in Music*, Academic Press (1992) 143–170.

Cosi, P., De Poli, G. & Lauzzana, G. Auditory modelling and self-organizing neural networks for timbre classification. *Journal of New Music Research*, v.23 (1994) 71–98.

Cope, D. *Computers and Musical Style*. Oxford University Press (1991).

Cross, Ian (1993). The Chinese music box. *Interface*, v.22 (1993) 165–172.

Davies, J.B. *The Psychology of Music*. Hutchinson (1978).

Desain, P. & Honing, H. The quantization problem: traditional and connectionist approaches. In M. Balaban, K. Ebcioglu, & O. Laske (eds.), *Understanding Music with AI: Perspectives on Music Cognition*, AAAI Press and MIT Press (1992) 448–463.

Goldman, C.V., Gang, D. & Rosenschein, J.S. NetNeg: a hybrid system architecture for composing polyphonic music. Proceedings of the *XI Colloquio di Informatica Musicale,* Bologna (8–11 November 1995) 3–6.

Kaipainen, M. Representing and remotivating musical processes: modelling a recurrent musical ecology. *Journal of New Music Research,* v.25 (1996) 150-178.

Kippen, J. & Bel, B. Modelling music with grammars: formal language representation in the Bol Processor. In A. Marsden & A. Pople (eds.) *Computer Representations and Models in Music,* Academic Press (1992) 207–238.

Kippen, J. & Bel, B. The identification and modelling of a percussion “language”, and the emergence of musical concepts in a machine-learning experimental set-up. *Computers and the Humanities,* v.23 (1989) 199–214.

Kunst, J. *Making Sense in Music: An Enquiry into the Formal Pragmatics of Art*. Communication and Cognition (1978).

Laske, O. Composition theory: an enrichment of music theory. *Interface*, v.18 (1989) 45–59.

Lee, C.S. The rhythmic interpretation of simple musical sequences: towards a perceptual model. In P. Howell, I. Cross & R. West (eds.) *Musical Structure and Cognition*, Academic Press (1985) 53–69.

Leman, M. Artificial neural networks in music research. In A. Marsden & A. Pople (eds.) *Computer Representation and Models in Music*, Academic Press (1992) 266–301.

Leman, M. *Music and Schema Theory - Cognitive Foundations of Systematic Musicology* Springer Verlag (1995).

Leman, M. Schema-based tone centre recognition of musical signals. *Journal of New Music Research*, v.23 (1994) 169-204.

Leman, M. Symbolic and subsymbolic descriptions of music. In G. Haus (ed.) *Music Processing*, Oxford University Press (1993) 119–164.

Little, D. Composing with chaos; applications of a new science for music. *Interface*, v.22 (1993) 23–51.

Lischka, C. On music making. Proceedings of *IX Colloquio di Informatica Musicale*, Genoa (13–16 November 1991) 80–84.

Longuet-Higgins, H.C. & Steedman, M.J. On interpreting Bach. *Machine Intelligence*, v.6 (1971) 221-41.

Longuet-Higgins, H.C. The perception of music. *Interdiscipilinary Science Reviews*, v.3 (1978) 148-156.

Sundberg, J., Friberg, A. & Frydén, L. Common secrets of musicians and listeners: An analysis-by-synthesis study of musical performance. In P. Howell, R. West & I. Cross (eds.) *Representing Musical Structure*, Academic Press (1991) 161–197.

Todd, N.P. A model of expressive timing in tonal music. *Music Perception*, v. 3 (1985) 33–58.

Widmer, G. Learning expressive performance: the structure-level approach. *Journal of New Music Research*, v.25 (1996) 179–203.

Widmer, G. The importance of basic musical knowledge for effective learning. In M. Balaban, K. Ebcioglu, & O. Laske (eds.) *Understanding Music with AI: Perspectives on Music Cognition*, AAAI Press and MIT Press (1992) 490–507.

Winograd, T. Linguistics and the computer analysis of tonal harmony. *Journal of Music Theory*, v.12 (1968) 2-49.