

## The systems approach to music

The Georgescus' visionary paper hews out a gigantic and magnificent scaffolding which encompasses an extraordinary breadth of musical phenomena. As is inevitable in a pioneering work of this magnitude, many of the details are elusive. Although a whole host of musical descriptions are given, the scope of the work is so ambitious as to permit limited space for individual details. In particular, it seems difficult to anchor the vision in a specific way to particular instances of musical development. Nevertheless, we find it more productive to concentrate on high-level aspects of the framework that is laid out than to carp about its lack of specificity.

**Characterizing the musical system.** The Georgescus regard the musical system as characterized by general properties such as wholeness, order, individualization, centralization, competition and finality. "Wholeness" refers to the fact that a musical work has an all-embracing integrity that transcends its individual components—as in any system, the whole is much more than the sum of its parts. "Order" recognizes that the work is nevertheless constituted from subsystems which operate at different levels within the organization of the whole, and these can be isolated and studied independently. "Centralization" refers to an integrative feature which pulls a work (or an entire genre of music) together, making all subsystems subservient to a single organizing principle such as a harmonic progression. The other terms are not defined in the paper, although their names hint at their meaning; they are presumably quite general properties which do not refer to specific musical events.

**Realizing these characteristics in a musical work.** These properties are realized through a vast number of different levels, or viewpoints, which ultimately result in a set of sound objects and processes. The levels, and their interrelationships, depend crucially on the historical and musical context of the composition. Layers include such attributes as melodic line, harmonic progression, tonicity, phrasing and use of silence, rhythm and rhythmic changes, dynamics, temporal and tonal transformations, and thematic structure. These layers are strongly interrelated and organized hierarchically. Through them the input to the compositional process—"the ensemble of human problems and meanings associated with a musical work, given as a structure in itself (neutral level)"—is transformed into the composition itself—"the musical work as an aesthetic, perceived reality." However, the Georgescus downplay performance issues, implying elsewhere that the output is musical *text*, uninterpreted by a human performer.

**Evolutionary development of music.** Few details or examples are given of how composers can use the range of musical levels to create works which express the characteristics they have in mind. Rather, the authors focus their main efforts on studying the evolution of the musical phenomenon from a system-theoretic standpoint. They see progress as a transition from the "flat wholeness" of primitive art to a "gradual differentiation of parts" through the birth of polyphony and harmony, the emergence of musical forms, tonal and themal contrasts, widespread stylistic diversification, and so on. Systemic concepts such as adaptation and self-organization may prove useful in analyzing this evolution. But the major insight is the interplay between periods of structural stability, punctuated by the turmoil of morphogenetic discontinuity. One is forcefully reminded of Kuhn's (1970) characterization of the scientific process as long periods of "normal science" interrupted by extraordinary, violent revolutions that question, dismantle, and reformulate the very foundations of the scientific edifice.

For example, the Georgescus point to the Classical structure erected by Hadyn and Mozart, and use the metaphor (or is it a model?) of the creative possibilities within a finite-dimensional state space being gradually depleted as the musical genre reaches maturity. At this point a Kuhnian paradigm shift occurs to a new, Romantic, style—created by Beethoven—with quite a different structure, to replace the functionally exhausted genre.

The natural system-theoretic model for a paradigm shift is a structural bifurcation or catastrophe. Periods of normal development within a particular genre are homeostatic: standing in dynamic equilibrium with the environment, adapting to change in an evolutionary manner, gradually exploring the possibilities afforded within this equilibrium. Eventually a tension develops between retaining and working within the current structural paradigm, and the creative quest for fresh avenues of variety and novelty without which emotional expression becomes stultified. The tension increases as the current possibilities become increasingly exhausted, until eventually a point of instability is reached and a revolution, or paradigm shift, occurs into a new structural form. A very clear example of the phenomenon can be seen in the Schoenbergian shift to atonality early this century.

**Instability within individual musical works.** The Georgescus daringly extend the vision of structural change from musicological analysis of stylistic evolution into particular compositions. This is a development that we find hard to assimilate. “Morphogenetic music” is defined as that which embodies structural change, as opposed to “normal” or structurally stable music. First, the Georgescus discard the idea of interpreting a particular musical work within the cultural context that created it, and instead regard the musical composition as a time-invariant dynamical system, “isolated in its perfection.” Second, they equate a musical work with the generative process that produced it. Whereas in certain genres—such as avant-garde stochastic and environmental music, or even traditional campanology—the work explicitly takes the form of a generative process, in most music a composition is the *result* of a generative process that takes place (or “executes”) in the composer’s mind. Finally, they contrast compositions produced by structurally stable processes with those whose processes contain instabilities. A changing-structure (morphogenetic) musical process is one whose structural instabilities outweigh the overall structural stability by some critical threshold. In such music, the alternate display of stable and unstable zones brings about violent changes in structure. While this definition is necessarily qualitative, the presence of structural instabilities is illustrated by some works of Beethoven, Bruckner, Shostakovic. In contrast, “normal” music is illustrated by works of Palestrina, Bach, Mozart, Beethoven, Schubert, Wagner, Ravel, Stockhausen, Messiaen and Kagel.

The Georgescus concede that the meaning of the concept of morphogenetic music is only “dimly seizable.” A relatively small number of good examples exist and they have hardly been investigated from this viewpoint. In a final section they elaborate on a particular kind of morphogenetic instability, the cusp-catastrophe model, and its relevance to music; but although the catastrophe model is simple enough (see Zeeman, 1976, for a particularly good introduction, and Poston & Stewart, 1978, for a more comprehensive treatment), we find its musical application somewhat unclear.

Our major difficulty with the identification of instabilities in particular pieces of music is that the mapping from work to process is extravagantly multi-valued. This is particularly true if works are analyzed in isolation from their cultural context, as the Georgescus seem to advocate. Within the theory of hermeneutics—the interpretation of texts—there is a continual tension between the idea that meaning lies solely within the text itself and the view that meaning is grounded in a process of understanding in which the text, its production, and its interpretation all play a vital part (see Winograd & Flores, 1986, for an introduction to hermeneutics from a modern cognitive perspective). Since

instability can only be a property of a process, not of a work in itself, attributing instabilities to a particular work seems to presuppose the identification of a process that explains or generates that work. In order to investigate stability, we need a mapping from work to process. Our objection is that there is no limit to the number of processes that generate a particular work. There is a danger that any instabilities are due to the particular mapping chosen, rather than being inherent in the musical composition. The problem is a fundamental one because no finite composition can unambiguously define a generative process.

## From musical surface to predictive model

The Georgescus' paper emphasizes the crucial role played by morphogenetic discontinuities within certain pivotal genres, and even individual compositions. The limitless structural possibilities afforded by such discontinuities has allowed the character of music to evolve through a panoply of widely differing styles. Nevertheless, despite the extraordinary richness offered by structural change, the task of characterizing "normal," structurally stable, music within a systematic framework is itself an ambitious and rewarding undertaking. Moreover, it seems to be a necessary precursor to the identification of morphogenetic discontinuities or structural instabilities, since these properties are possessed by processes, not individual works. The analysis required to derive a generative process from a particular composition must be driven by a theory of normal, structurally stable, music. Only then will it be possible to identify and exploit the richness and intricacies offered by morphogenesis.

In the Knowledge Science Institute of the University of Calgary we have underway a research program aimed at identifying structure in widely differing genres of music. The purpose is to explore new methods of predicting music, and ways of evaluating the effectiveness of prediction methods. Like the Georgescus, we focus on the *musical surface* of a piece—that is, the notes, their durations, and assignment to voices—and neglect performance issues. Prediction is based on combinations of structural, mathematical, and statistical methods. Evaluation is based on "compression," a strong measure of the predictive power of a probabilistic model.

Our methodology is highly empirical in nature. Predictive models are derived automatically from computer analysis of large bodies of existing music. We are beginning with a detailed study of the Chorales of J.S. Bach. The predictive power of our models will be compared with the that of human experimental subjects. The goal of this project, boldly stated, is to construct a system that rivals the performance of the best human subject in guessing the musical surface of a Bach Chorale, given incomplete information about that Chorale. Whether this is possible using any existing paradigm for music prediction is quite unknown at this time. Regardless of the outcome, strong statements can be made regarding the types of context, knowledge, and background used by a musician. In effect we are addressing the issue of how effectively a purely syntactic system can model the (structurally stable) musical surface.

The relevance of our endeavor to the Georgescus' is that only through automatic identification of predictive models of particular pieces and genres can morphogenetic instabilities be investigated and analyzed in a fully objective manner. We share the view of Meyer (Meyer 1956) that instability results from temporarily or permanently blocked expectations on the part of the listener. These points of deviation are of interest in identifying important structural points in the music. In order to detect such instabilities, a model must be formed which creates expectations, or "predictions," in order that their breakdown may be detected. The aim of our research program is to create such models.

The following sections briefly survey the musical background to our quest, set out the methods being used to create and evaluate predictive models, and sketch the design of an experiment designed to investigate human performance in predicting music.

## Contributions from musicology

**Schenkerian analysis.** The foremost example of a purely structural theory for music is that of Heinrich Schenker (1868–1935) (Schenker, 1935; see also Forte, 1959, and Jonas, 1934, for good introductions). Schenker views music, through a series of diminutions of the surface, as comprising three basic levels. The background level is based on one of eight basic “prototypes,” each of which describes a descending melodic line over a I–V–I bass arpeggiation. While it is far beyond the scope of this note to debate Schenker’s theory in detail, many aspects make it computationally intractable and in some respects inadequate for our work.

First, the theory does not explicitly state how one level can be rewritten, or transformed, to another. Indeed one rarely encounters two identical Schenkerian analyses of a given composition. Computer implementations invariably rely on a few simple, well-defined transformations. For example, Frankel *et al*’s (1978) “PT” transformation fills in gaps between adjacent notes, while the “UA” function inserts auxiliary tones between identical adjoining notes. The approach is not only incomplete but stratifies a series of transformations into a sequence of levels, emasculating the rich interplay between them that Schenker envisaged. Second, no axioms are provided for predicting durational values on the musical surface. Given a background level, Schenker does not specify durational values received by notes inserted at lower levels. There is no requirement that background events follow well-defined dispersions throughout a score: one often sees analyses in which a large proportion of background notes occur in strict temporal contiguity on the musical surface. The final shortcoming, and perhaps the most important, is the heavy reliance that is placed on the theoretician’s tacit knowledge. It is quite unclear what that knowledge is and how it contributes to the analysis. To render it explicit enough to be usable for prediction would constitute an ambitious research project in musicology.

**Structural and grammatical approaches.** Lerdahl and Jackendoff (1983) describe an impressive and comprehensive application of hierarchy to music whose most notable feature is the modeling of meter. A metrical analysis combines with a grouping analysis to form a “time-span reduction.” Such a reduction can progress over multiple levels—there is no strict categorization of levels as in Schenker’s theory—and is accomplished by two types of rule: well-formedness and preference rules. The former dictate the bounds, limitations, and structure of analysis. For example, one such rule ensures that subordinate groups of notes in a hierarchy occupy the same higher-level group. The latter constrain analysis by allowing only reductions corresponding to an experienced listener’s intuition. For example, one such rule prefers analyses that “most closely approach the ideal subdivision of groups into two parts of equal length.” Despite its attractive ability to link metrical hierarchy with pitch hierarchy, it is not easy to see how the theory can be used generatively. Although the analyses presented are multivoice, predictions are only offered for the main melody line.

Some researchers (eg Roads, 1979; Cope, 1987; Baroni & Jacobini, 1978) have taken the “music as language” metaphor more seriously, attempting to apply various formal grammars directly in a generative fashion. Recent research, however, has made it apparent that grammars for music may be quite different from grammars for natural language (Blacking, 1984).

**Mathematical characterizations.** Structural theories permit the description of groups, structural boundaries, and the broad features of a work. When it comes to making a prediction, however, current theories are insufficiently explicit. On the other side of the spectrum are numerical methods which have tended to concentrate on local attribute relationships, ignoring form, structure, and repetition (eg Hiller & Isaacson, 1959). Generally statistical, such descriptions of music are often criticized as superficial. Meyer (1957) offers a relatively positive high-level discussion of the operation of entropy, information, and Markov processes in music.

Lewin (1987) describes a wide-ranging attempt to formalize relationships between elements in a musical space using various metrics. Questions of the form “if I am at  $s$ , and wish to get to  $t$ , what operation must be performed on  $s$ ?” are considered, where elements  $s$  and  $t$  can be any musical gestures of the same type. The ideas are first developed using mathematical group theory, although later a more general approach using transformational graphs or networks is taken. This theory seems to offer a good basis for predicting the musical surface. On closer examination, however, it is much too strict. How can an “interval” be defined between two phrases in a Bach chorale? It is quite unrealistic to seek a mapping which when applied to the first phrase yields the second, as advocated by the theory. A more workable solution would be to provide mappings between the main structural events of the phrases. One can imagine organizing these into a hierarchy resembling those of Lerdahl & Jackendoff (1983), with mappings at each level. A direct-product interval system between melodic and rhythmic spaces could avoid the isolation of either parameter, providing an analogue to the linked time-span reductions discussed above.

Baffioni *et al* (1981) adopt a statistical approach with a strong structural flavor. They view music as evolving on different time-scales. Important structural events in a sequence have a high entropy. Many time-scales may co-exist: one spanning the duration of the work; others at the level of section, phrase, motif; all the way down to individual notes. Each level is only effective as a predictor at intervals of its own time-scale.

## **Adaptive modeling of music**

Our methodology derives from extensive work on adaptive modeling and prediction of text (eg Witten & Cleary, 1986; Bell *et al*, in press). The technique we use to evaluate models is to measure the entropy, or information content, of compositions in the genre with respect to each model. This is basically the amount of extra information that is needed to specify the composition, given the model. For example, using a model that captures a music student’s knowledge of the form of Bach Chorales, together with the usual rules of melodic development, harmonization, phrasing and rhythmic gestures, a particular Chorale can be communicated using much less information than would be needed in the absence of the model.

The techniques we use create *probabilistic* models that assign specific probabilities to the selection of each symbol in context. Then the entropy of a particular symbol sequence can easily be evaluated. Moreover, the entropy has an operational interpretation as the minimum number of bits (binary digits) necessary to communicate the sequence, given the model—a result due to Shannon (1948). Furthermore, and even more surprisingly, there is a fairly simple algorithm to encode the sequence into just this number of bits, and to decode the resulting bit stream to recover the original sequence (Witten *et al*, 1987).

**Creating models of sequences.** Work in adaptive text compression explores how much of the surface structure of text can be modeled without resorting to a predetermined grammar. Written text is viewed as a purely linear syntactic system. Probabilistic finite-context models are formed from a corpus of written text, and are used to predict (or, equivalently, “compress”) further entries in the corpus. The success of the compression is measured in terms of the average entropy of predictions. Techniques for text compression operate on symbol (letter or word) sequences in strict sequential order. Temporally ordered models can be derived from a source simply, directly, and elegantly. They can be used to compress or to regenerate the symbol sequence. Substantial difficulties of grammatical inference arise when more powerful languages are used (see Gold, 1967; Angluin & Smith, 1983).

**Creating models of multi-level systems.** It would be extraordinarily simplistic to regard music as a one-level system. Parallel streams of expression render it at the very least an entwining of several sequences of events. A natural way to extend techniques for modeling temporally-ordered symbol sequences is to adopt multiple viewpoints. A “viewpoint” is a knowledge source that regards a stream of events in a particular way, allowing predictions to be made about surface events. It is logical to expect predictions to improve when various sources of knowledge about an event stream are correlated and used together. Different viewpoints predict at different time-scales. They need not predict just the next event in the sequence, but can predict events which occur at any point in the future. Thus viewpoints can naturally and effectively accommodate multiple parallel streams of activity and arbitrary partial temporal orderings—much richer possibilities than those afforded by phrase structure grammars. The use of multiple viewpoints has precedents in Andreae’s (1977) “threaded contexts” and Witten & Cleary’s (1983) adaptive modeling of quadrees.

**Modeling pitch and rhythm.** Pitch movement provides a nice example of the utility of multiple viewpoints. Relevant viewpoints include

- interval from the tonic, that is, the distance of a note from the tonic of the piece;
- interval from the preceding note in the same voice;
- contour with respect to the preceding note—rising, falling, or static.

(Lewin (1987) outlines a technique for the construction of new intervallic pitch viewpoints.) A system that models pitch in these terms, incrementally and adaptively, has been designed and tested on several Gregorian chants with promising results (Conklin & Cleary, 1988).

To predict the complete musical surface, it is necessary to model and predict the durations of individual notes. Alternatives are

- to provide explicit rhythmic viewpoints;
- for each element of a pitch viewpoint, to attach a time-scale at which the prediction applies.

The problem with with the former is that rhythmic information is ignored when modeling pitch (and *vice versa*). The problem with the latter is that durations are too tightly bound to specific pitch movements, providing too fine a granularity for flexible modeling. The underlying problem is that on the one hand durational information cannot be separated from the pitch viewpoints, while on the other parameters must not be linked too closely. A good model is a *creative* one—a model that is flexible enough to apply in diverse situations will provide better compression than a narrower one.

**Multiple voices.** The problems associated with modeling multivoice music, homophonic or polyphonic, will probably prove to be the most challenging and difficult part of this research program. In earlier work a *vertical interval* viewpoint was constructed so that if the vertical interval from a piece matched the context of a rule, a new vertical interval could be predicted (Conklin & Cleary, 1988). There is a strong case for modeling vertical intervals separately for each pair of voices. For example, a perfect fourth interval between bass and soprano voices is quite different to the same interval between tenor and alto voices. For a four-voice composition this requires six different viewpoints—an unacceptable level of redundancy. Instead, a “vertical block” viewpoint can be adopted, each of whose elements contains a component for every voice. Components could be any of the pitch viewpoint elements discussed above. A problem analogous to that discussed above for rhythm is raised: one method is too weak, omitting information about voice-specific treatment; while the other is too rigid, predicting strictly for *n*-voice works, and overlooking similarities between different pairs of voices.

**Harmony, and figured bass.** Obtaining a good algorithm for parsing tonal harmony is a problem fraught with difficulties (see Winograd, 1968, for a parser for tonal harmonic movement based on a systemic grammar; Hindemith, 1937, proposes a detailed method). The problem is a fundamental one: to construct a theory that balances and encompasses both homophonic and polyphonic music. If a composition were annotated with figured bass, this information could be used by creating a new viewpoint for it. Note, however, that the interpretation of a figured bass prediction is by no means well defined. How would it instantiate the alto voice? Certainly the probability of a chord tone is high, but an ornamental tone is also likely. Thus such a viewpoint alone does not solve the prediction problem.

**Creating new levels.** A mechanism is necessary to identify crucial events in a piece so that they can be given a more important structural role. A technique derived from Baffioni *et al* (1981) could extract events having high entropy with respect to some viewpoint and promote them to the next higher structural level. Each level would be modeled in the same way. At the top level, only the highest entropy event for the viewpoint remains. On generation, this would be predicted first by the viewpoint.

Entropy with respect to viewpoint is not the only cue to structure. Other indicators are diminished rhythmic movement, repetition of viewpoint elements, and phrase boundary indicators visible on the score. We are exploring the extent to which a measure of entropy can encompass techniques like these.

## Human prediction of the musical surface

How can a model of music be evaluated? The obvious way is to generate sample compositions from it and gauge their quality. This suffers from the drawback that assessment can only be subjective. A second possibility is to examine the model itself to see the extent to which it incorporates the musical rules that govern compositions in the genre. This is a very tedious undertaking since models formed adaptively are generally large and opaque. Even if it were possible to extract the rules embodied in the model, it would often be very difficult to assess whether or not they reflected genuine musical principles. Since we seek a methodology that is capable of capturing regularities in different musical genres, any evaluation must be relative to the particular style being modeled. Considering the extensive and widespread knowledge of the form of a structurally stable genre such as the Bach chorales, it may be relatively easy to evaluate rules formed for this genre. However, problems arise when one attempts to assess rules that govern compositions in other genres, say bebop improvisations in the style of Charlie Parker, or Indian ragas. A third evaluation technique is to measure the entropy, or information content, of compositions in the genre with respect to the model. This is an excellent yardstick when the performance of various viewpoints and time-scale configurations must be compared and contrasted, but does not relate to human abilities or cognitive processes. It would be nice to be able to compare automatically-generated models for a particular genre of music with those of human musicians.

We are undertaking an experiment which will gather data on human performance in constructing and predicting Bach Chorales. Similar experiments have been performed for written text (Cover & King, 1978) following pioneering research by Shannon (1951). Subjects are presented with an incomplete fragment of a Bach Chorale on a computer screen, and asked to guess the musical surface, note by note. They are given the full rhythmic skeleton in advance. The subject points to a position in the skeleton and enters a guess for the note at that position. The note's correct identity is not revealed until it is guessed correctly. The subject need not work from left to right but may move freely around the score.

The data collected comprises (position, guess) pairs which partially capture the subject's plans, goals, and strategies. The average entropy of the piece with respect to the subject's model can easily be calculated. The computer will be given the same tasks, and its performance compared to that of human subjects.

## Conclusion

This paper has taken a wide-ranging tour of the systems approach to music. The Georgescus identify the system-theoretic notion of a structural bifurcation with a paradigm shift in musical expression. Then they go further and speak of identifying such instabilities within individual musical works. In order to follow this line of reasoning we feel compelled to model individual musical works by a dynamic, predictive model—for it seems to us that morphogenesis can only be associated with a creative process, not an artistic product *per se*. Structural discontinuities correspond to a breakdown of the listener's predictive model of the music.

This leads naturally into our own research on deriving predictive models of musical works adaptively from the musical surface. We equate predictive power with the ability of a model to compress the surface of unknown pieces. An examination of now-standard musical theories such as Schenker's (1935) analytic framework, and more recent approaches such as Lerdahl & Jackendoff's (1983) generative theory and Lewin's (1987) notion of generalized musical intervals and transformations, by no means solve the problem of forming predictive models. On the other hand, work on adaptive modeling and compression of text—while it provides an excellent yardstick for measuring predictive power in terms of entropy—is limited to a single stream of events in strict sequential order.

Our solution is to adopt a multi-faceted theory of prediction, where a variety of viewpoints (eg diatonic, chromatic, relative interval, for pitch) combine to predict aspects of the musical work at various time-scales. This formulation allows us immediately to confront the problem of how different viewpoints—such as rhythmic and melodic ones, or the relationship between different voices—are linked. While we have not yet solved these problems satisfactorily, at least we can identify them and evaluate alternative ways of tackling them. Finally, there is a need to compare the performance of our adaptive models with musicians. We have shown how a benchmark level of performance can be established by testing human subjects on their ability to predict the musical surface of examples taken from a restricted genre.

Once we can reliably establish predictive models from the musical surface, we will then be in a position to continue the Georgescus' ambitious research program, identifying structural instability in musical works by seeking unfulfilled predictions.

## Acknowledgements

We would like to thank John Cleary, Brent Krawchuk, Leonard Manzara and Dave Mauksby for useful discussions about music modeling. This research is supported by the Natural Sciences and Engineering Research Council of Canada.

## References

- Andrae, J.H. (1977) *Thinking with the teachable machine*. Academic Press, London, England.
- Angluin, D. and Smith, C.H. (1983) "Inductive inference: theory and methods" *Computing Surveys*, 15 (3) 237-269, September.
- Baffioni, C., Guerra, F., and Lalli, L. (1981) "Music and aleatory processes" *Proc 5-Tage-Kurs of the USP Mathematisierung*, 273-303, Bielefeld University, Germany.
- Baroni, M. and Jacobini, C. (1978) *Proposal for a grammar of melody*. Les Presses de l'Universite de Montreal, Montreal, Canada.

- Bell, T.C., Cleary, J.G., and Witten, I.H. (in press) *Text compression*. Prentice Hall, Englewood Cliffs, NJ.
- Blacking, J. (1984) "What languages do musical grammars describe?" in *Musical grammars and computer analysis*, edited by M.Baroni and L.Callegari. Leo S. Olschki, Florence, Italy.
- Conklin, D. and Cleary, J.G (1988) "Modelling and generating music using multiple viewpoints" *Proc First Workshop on AI and Music*, 125-137, American Association for Artificial Intelligence.
- Cope, D. (1987) "An expert system for computer-assisted composition" *Computer Music Journal*, 11 (4) 30-46.
- Cover, T.M. and King, R.C. (1978) "A convergent gambling estimate of the entropy of English" *IEEE Trans Information Theory*, IT-24 (4) 413-421, July.
- Forte, A. (1959) "Schenker's conception of musical structure" *J Music Theory*, 3 (1) 1-30.
- Frankel, R.E., Rosenschein, S.J., and Smoliar, S.W. (1978) "Schenker's theory of tonal music—its explication through computational processes" *Int J Man-Machine Studies*, 10 (1) 121-138, January.
- Georgescu, C. and Georgescu, M. (1989) "A systems approach to music" Unpublished research report.
- Gold, E.M. (1967) "Language identification in the limit" *Information and Control*, 10, 447-474.
- Hiller, L. and Isaacson, L.M. (1959) *Experimental music*. McGraw Hill, New York.
- Hindemith, P. (1937) *The craft of musical composition: Book 1, Theoretical Part*. Schott and Co.
- Jonas, O. (1934) "An introduction to the theory of Heinrich Schenker" Translated and edited by J. Rothgeb, Longmans, New York.
- Kuhn, T.S. (1970) *The structure of scientific revolutions*. University of Chicago Press, Chicago, ILL.
- Lerdahl, F. and Jackendoff, R. (1983) *A generative theory of tonal music*. MIT Press, Cambridge, Massachusetts.
- Lewin, D. (1987) *Generalized musical intervals and transformations*. Yale University Press, New Haven, CO.
- Meyer, L.B. (1956) *Emotion and Meaning in Music*. University of Chicago Press, Chicago, ILL.
- Meyer, L.B. (1957) "Meaning in music and information theory" *J Aesthetics and Art Criticism*, 15, 412-424.
- Poston, T. and Stewart, I. (1978) *Catastrophe theory and its applications*. Pitman, London, England.

- Roads, C. (1979) "Grammars as representations for music" *Computer Music J*, 3 (1) 48-55.
- Schenker, H. (1935) *Free composition (Der freie Satz)*. Universal Edition A.G., Translated and edited by E. Oster, Longmans, New York, 1979.
- Shannon, C.E. (1948) "A mathematical theory of communication" *Bell System Technical J*, 27, 398-403, July.
- Shannon, C.E. (1951) "Prediction and entropy of printed English" *Bell System Technical J*, 50-64, January.
- Winograd, T. (1968) "Linguistics and the computer analysis of tonal harmony" *J Music Theory*, 12, 2-49.
- Winograd, T. and Flores, F. (1986) *Understanding computers and cognition*. Ablex, Norwood, NJ.
- Witten, I.H. and Cleary, J. (1983) "Picture coding and transmission using adaptive modelling of quad trees" *Proc International Electrical, Electronics Conference*, 1, 222-225, Toronto, ON, September 26-28.
- Witten, I.H. and Cleary, J.G. (1986) "Foretelling the future by adaptive modeling" *Abacus*, 3 (3) 16-36, Spring.
- Witten, I.H., Neal, R., and Cleary, J.G. (1987) "Arithmetic coding for data compression" *Communications of the Association for Computing Machinery*, 30 (6) 520-540, June.
- Zeeman, E.C. (1976) "Catastrophe theory" *Scientific American*, 234, 65-83.