Music In The Air: a theoretical model and software system for music analysis and composition

ROBERT SPALDING NEWCOMB

Director of Information Technology, University of Michigan School of Music, Ann Arbor, MI 48109-2085, USA
E-mail: rnewcomb@umich.edu

As a composer in the very late twentieth century my goal is to inject as much intelligence and substance into the process of music-making as possible, to present listeners with sensory and intellectual metaphorical tools that enable them to perceive relationships of time and space in ways that might otherwise go unnoticed. Music-oriented researchers in the fields of psychoacoustics and cognition in recent years have viewed the computer as a means of simply extending and rationalising traditional music theory and form, rather than summarising the base of musical knowledge we have inherited and then breaking free of those constraints. Performers, composers and producers in the field of popular entertainment continue to use technology to narrow rather than broaden aesthetic boundaries in the so-called creative arts.

A new aesthetic paradigm is needed . . .

Technology provides us with the tools we require in order to build a genuinely new method of music creation and listening. I propose that it is time to invest our energies in the creation of visionary learning tools that enhance our ability to evolve in our relationship to ourselves, to one another, to the planet, and perhaps ultimately to the stars. It is the author's strongest desire to allow for the full realisation of what computer music has long held the promise to provide – limitless variation with infinite control.

Toward this end, I document here (i) a theory of music that attempts to assimilate the infinite and subtle interconnectedness of the human experience into a single symbolic structure, and (ii) the application of this theory to the extension of an emerging genre – the environmentally interactive computer music system – designed for eventual use by people not necessarily aware of the underlying technical and aesthetic principles at work, but fully appreciative of the profound personal benefits to be found embedded therein.

Owing to space limitations, I will focus almost exclusively on the theoretical and aesthetic issues that have driven the software development, and leave a detailed discussion of the system specifications and its operation to a separate article.

1. INTRODUCTION

Music In The Air (MITA) v0.71 is a computer-based music composition system, a sophisticated applications environment capable of realising offline and realtime musical work of singular originality and quality. The core of the design is based on establishing integral links between improvisational musical performance, algorithmic composition, natural-world non-musical phenomena, and a user-interface that effectively monitors and guides the interaction of these high-level event streams.

Compositions are produced by following a two-step process (figure 1). Musical material, or alternatively non-musical material of any origin, is input to a preprocessor classification system that analyses the mathematical structures embedded within the raw data. The analysis is based on a user-customised set of attribute definitions. The information gained from this analysis is then used by an assembly process to construct original musical structures that reflect the structures identified in the input data. In addition to the knowledge derived from the source material, this process is informed by several self-perpetuating control mechanisms that interact with one another and with a user for the duration of assembly.

The system includes an intentionally critical dependence upon 'realtime' user and environmental input during the actual compositional construction process. The design is influenced by the physical and biological sciences through the theories of universality, chaos (sensitivity to initial conditions), and genetics among others, and by mathematics through techniques such as statistical determinacy, probability, distribution, classification, and hierarchical transformation processes.

The aesthetic battleground from which the system has emerged is framed by the duality of 'response-time within a performance context' versus 'depth-of-intelligence and scope-of-prior-knowledge within a nonrealtime compositional context'.

2. BACKGROUND

Generally, attempts to explain the processes of musical composition, interpretation and appreciation have assumed that human emotional inspiration, a need to communicate with others, a focus of will upon available musical materials, and an immense storehouse of auditory memories form the basis for study. Another school of thought claims these concerns to be too subjective, and has become committed to a process of music creation based exclusively on statistical

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

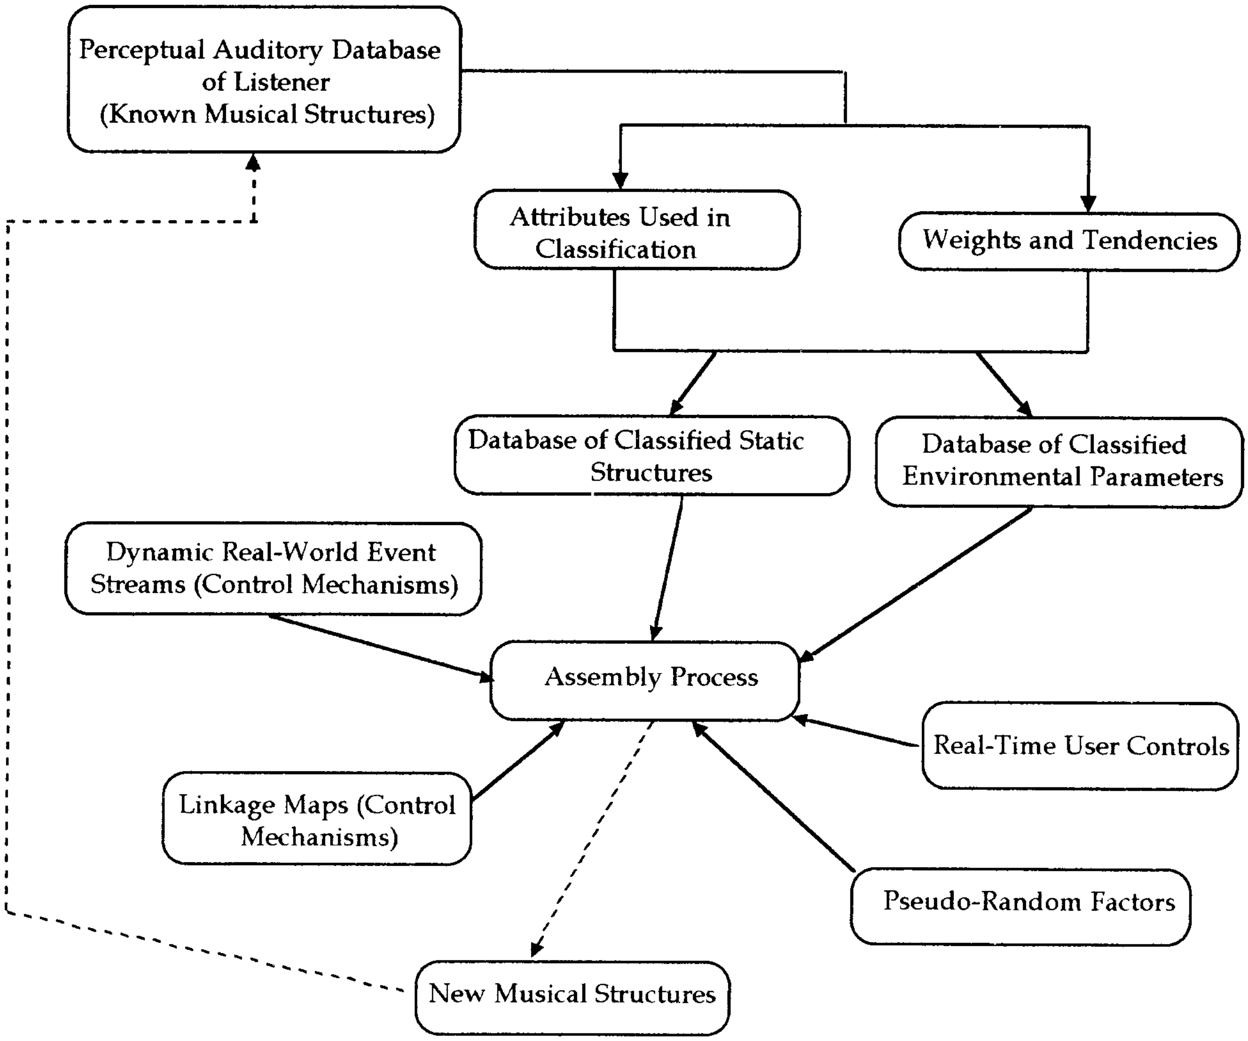

Figure 1. General model of cognitive musical structure.

algorithms that do not reference emotional sources. Others have modelled compositional software after processes found to be essential to biological and biophysical phenomena, or base their research upon an analysis of the processes found in improvisation.

Separately, each of these perspectives fails to measure the entirety of the phenomenon – it is all and more than the sum of these components. While we do control much of what we produce as composers and what we in turn consume as listeners, there are in fact many critical factors left unexplained in distinguishing brilliant from mediocre, vital from boring. As sentient context-bound entities, humans do in fact carry a vast depth of preconditioning to any listening experience, thereby implicitly validating much 'culture-based' research. Statistical methods originated in the study of natural phenomena and bring tremendous analytical power to the study of process and structure. The analysis of improvisational performance has taught us much about how we depend upon the manipulation of memorised materials for our spontaneous compositions.

Scientific research during the last three decades has led to intriguing hypotheses concerning the nature of humankind's relationship to the terrestrial biosphere and to the world beyond this planet. Music researchers have been involved for years in breaking down barriers between naturally occurring processes and the artificial, culturally imposed methods involved in traditional music composition. The importance of natural processes in any human activity cannot be denied, and their simultaneous simplicity and complexity continually yield revelations with regard to the nature of 'creativity' and form.

In designing computer-assisted composition systems, researchers are immediately confronted with the issue of prioritising the relative importance of realtime system performance versus the depth and complexity of the musical environment upon which the system will operate. What do we mean by the

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

terms 'realtime', 'interactive', 'significant' and 'intelligent', and how is a relative value placed upon phenomena displaying these characteristics?

Perhaps we must confront an even more basic challenge: How does one integrate into a single *conceptual theory* and *practical system* the vast territory of *reductive* analytic methods developed in the tradition of Music Theory, with *constructive* techniques using the plethora of technological materials and tools available today? By failing to form a theory of music that consistently integrates these two processes, we have promoted the social/cultural fragmentation that supports the *destructive* mythology of 'creators versus performers versus critics/theoreticians' in musical history. If the fruits of rigorous analysis do not explicitly and directly form the basis of subsequent compositional work, which in turn is evaluated based on the system from which it emerged, then we have accomplished nothing other than furthering the incessant indulgence in distraction so prevalent in artistic and musical activity.

In order to adequately address what the author feels to be the most important characteristics of musical expression, the system described here is an attempt to build firstly upon structural foundations that simply cannot presently be dealt with sufficiently in real time. These foundations are based largely upon the experience of improvisational expression, human memory, and upon an observance of Nature's systems of operation.

In a traditional offline compositional process, composer and computer interact for 'as long as it takes to create a significant piece of music'. In support of this approach, several software developers have opted to concentrate on extending the familiar traditional paradigm of offline 'create and edit' processing. Others follow the dictum of interactive realtime performance that states that 'we must write or invent as fast as technology allows', implicitly limiting one's ideas to those that are executable within the bounds of acceptable definitions of 'realtime'. The majority of this type of effort has so far been directed toward realtime performance systems that have relied heavily upon simple trigger mechanisms for their musical decision-making.

In the MITA system, I have attempted to provide a more sophisticated link between Nature and composed sound, not building what tend to become very shallow cause-and-effect couplings, but examining data from these naturally occurring event streams in order to learn of their inner workings. Translation mappings, to be effective, must do justice to the complexity of both the source information and the resulting musical form.

The concept of interactivity is interpreted somewhat differently than in other existing systems. The time required to create a composition becomes indefinite, arrived at mutually by composer and the system's control mechanisms, so that even though a realtime performance is not necessarily taking place, the process of composition is effectively a 'live, realtime compositional process', propelled by its own momentum, continually in a productive state of development, regardless of the composer's level of activity. Effectively, MITA exists in the vast unexplored territory lying between pure 'realtime' and pure 'offline' methodologies.

3. THEORY

3.1. Music as metaphor

Music composition is a powerful metaphor for any number of global problems facing the human race. The best solutions to nearly all of our political, economic, sociocultural and interpersonal problems reduce, given enough analysis, to a strikingly similar conceptual paradigm: the development of a system that recognises, accepts and nourishes innate differences, yet accomplishes this by adhering to a vision that is also profoundly committed to balancing the relative values of each differing component, dynamically redefining itself to accommodate new strata of differing perspectives as they appear, and handling 'transactions' between its elements in a consistently unbiased manner.

3.2. Auditory databases

A major issue that evolves from this perspective is the concept of familiarity in the human auditory experience. Each human mind–body complex is unique, and therefore maintains a unique database of listening experiences. This database represents the receiving matrix to which a composition transmits ideas. An analysis of the listener's database of musical memories produces a list of behavioural tendencies that is difficult but not impossible to quantify. One of the many characteristics found will undoubtedly be the relative adventurousness that exists in the listener. The most adventurous will seek out compositions that reflect the most divergence from the experiences stored in her musical memory – forms that do not adhere to stylistic domains known to the listener are perceived as desirable, valuable data. The least adventurous listener will seek out compositions that are similar to those in her memory that she has enjoyed. She will be unable to handle information that lies very far outside her cumulative understanding of music, therefore missing much of what is occurring within the composition.

3.3. The arbitrary and synthetic nature of sequential logic

No decision is any more arbitrary than any other decision. However minute or basic, all decisions lead to some consequence. All decisions move a composition, a career, a musical era, or a planetary epoch, forward into becoming what only it can be – the unique result of a massive number of decisions, based on an even more massive number of criteria, or system constraints.

This is not an insignificant observation. It is in fact the basis for the general model of music composition described in this paper, a summary of a concept upon which I have been building for years, 'the arbitrary and synthetic nature of sequential logic' (Newcomb 1980, 1981). The challenge at hand is to determine what decisions affect what other decisions and to what degree, relative to all other relevant decisions.

It must be understood and accepted that any sequence of decisions followed in building a temporal structure (quantifiable data that vary within a time domain) is a valid process, maintaining its unique definition of consistent, rational logic. Therefore, the resulting constructions must also be considered worthy of inspection, and inevitably, of some degree of acceptance or rejection. Of utmost importance is that we recognise these musical compositions as residual evidence of the altogether new 'languages' that created them. Though a system has been initially designed by humans, the languages and the compositions it subsequently creates are quite literally intelligent and conscious as they are growing. We as designers can learn immeasurably from them. This must be the prime motivation for building a system in the first place, to learn from the data that it eventually produces. These compositional structures can only be understood if evaluated on their own terms, not on the terms of pre-existing vocabularies of decision-making hierarchies.

3.4. Extended symmetry as a basis for evaluating structural integrity

The concept of relative 'symmetry' is often a primary point of departure in evaluating the integrity of a structural model. The generally accepted definition of the word needs to be extended if it is to serve adequately for the purpose of addressing my concerns in constructing a general model of the compositional, and in fact, the creative, process.

Symmetry must be understood to be a measurement of design and production consistency existing within and between all scales of model attributes, from minutiae to macrocosm. The system of thought responsible for the model must account for structural relationships within each scale, as well as for all possible relationships existing between each scale and each other scale.

Given this extension of interpretation, the analysis of symmetrical characteristics in any given structure produced from a general model must therefore be based upon a thorough understanding of the design of the general model itself. If the design of the model contains attributes that sufficiently reflect symmetry, then it follows that each individual structure produced from the model implicitly retains this structural integrity. The challenge in recognising symmetry in a single case lies in maintaining an awareness and acceptance of the inherent variability of how the symmetry within a general model becomes manifest within each of its progeny.

3.5. At issue: cultural and personal norms in need of realignment

The arts and sciences have evolved into convergent cultural event streams. The primary function of these two disciplinary arenas of exploration must be to broaden the perspective of cultural and personal awareness, and to regenerate a healthy balance within individuals and cultural institutions.

3.6. Computer-generated music composition – a paradigm shift long overdue

In setting forth guidelines for the realisation of computer-generated compositions that are in essence guaranteed to be not only 'original' but 'good' music, an implicit rebellion is also being described. The crucial point to realise is that a rebellion is not a rejection. All of the work that I, and others, have done in promoting serious consideration of the consequences to be found in musical applications of self-organising systems, self-determining rule systems, and intelligent knowledge database analysis and construction methods, is an extension of historical musical work. None of our work is possible without a coherent understanding of the systemic limitations we have inherited. Many traditional theorists, composers, musicians and educators, who scoff at the thought of music 'made by machine', in their contention that 'we have gone too far', belie their fear that 'they have not gone far enough' and cannot personally or professionally cope with the philosophical implications of the furthest reaches of machine intelligence. To these supporters of the musical foundations upon which I build continually, I present the explanation contained herein, in hopes of allaying any fears, and to offer assurance that the foundations are secure, and may be left unattended as the ongoing design of contemporary musical architecture is developed.

The paradigm shift referred to has come about due in large part to the dissemination of massive doses of musical information in the twentieth century. The acceleration of this outpouring of musical expression shows no sign of slowing. It does, however, make clear that the vast majority of musical expression is in fact focused on the expression of an extremely minute fraction of what the human race is capable of expressing, and in fact must express in order to adequately recognise the nature of our global situation.

Musical expression is a means of encapsulating emotional, psychological, physical and intellectual experience in an abstract form, eventually resident in the temporal domain of a stream of varying air pressures. This stream is perceived by us on all the levels through which it was created, emotional, psychological, physical and intellectual, thereby ensuring a channel of communication between two or more human beings. The communications protocol for this channel is where the rebellion must take place. We are a population that has embraced everyday technological 'communications' advances such as automatic teller machines, cable television, cable-less television and wireless telephones. Yet, musically speaking, we are using tremendous advances in computer and audio technology to basically communicate at the level of sitting around a fire in a cave, banging two rocks together.

The constraints that entrap popular music today have entrapped it since 'popular' music was first identified as something distinct. The same is true of what we call 'serious' music, of 'rock' music, of 'folk' music, of 'gospel' music, of 'blues' music, 'rhythm and blues' music, 'jazz' music. All popular music sounds the same. It sounds like popular music. All jazz music sounds the same, etc. All of these 'styles' of music have their nice little set of rules by which all music must play if it is to be nominated for inclusion in a given style. These rule sets deal not only with structural characteristics but with characteristics of emotional expression as well, in essence dictating what information is allowed to pass through our communications channel from one human to another using our musical protocol. So, an individual has a choice of what style she will 'customise' (within the rules though), to use as her vehicle of musical expression. Individuals do not define styles. The fate of an individual whose work does not fit a style's definition will be either (i) not to be heard at all (censure), or if enough potential commercial and/or professional gain for all other purveyors of the style seems likely, then (ii) to be absorbed into a 'style' with the understanding that the individual has 'expanded' the scope of the style's 'territory', and therefore has acted as an ally on behalf of the 'style rules committee'. This scenario seems to completely eliminate the possibility of expressing anything remotely original, reflecting an individual's unique perception and experience of existence. In political terms we call this a martial state, facism. In socio-cultural terms we call it paranoia, homophobia. Ironically, and quite cleverly, champions of stagnate traditions are themselves celebrated as preservationists, purportedly providing a cultural service, and therefore receive the vast majority of private and governmental financial support. This obviously creates an even more hostile environment for original thinkers working in music composition.

All rule sets are made to be expanded upon, granted, but the concept of a pre-existing rule set itself is where our whole musical communications system has failed us. The protocol is poorly designed, and must be overhauled from the bottom up, rethinking all assumptions that have been made. In fact, there can be no assumptions made at all.

A rule set that evolves over time and is based on any sort of consensus or agreement between several individuals, communities, or even institutions, represents a wealth of stability. But it also creates caricature, stagnation, and encourages a dangerously subversive acceptance of artistic limitations. Any rule set will attract inherently insecure individuals as proponents and admirers, who will invariably become either entrapped within its confines or move on to another rule set as they seek a platform for musical experience.

MITA attempts to defeat all existing and potential rule sets by providing a model of the very process of identifying rule set traits, collating them, and then proceeding to obliterate the boundaries between them during the creation of new structures.

Differentiating between the tradition of the empirical 'cut and paste' approach to artistic creation and the recently evolving tradition of self-organising structures is essential for an appreciation of MITA and of the inherent constraints and freedoms existing in both conceptual frameworks. The most basic element to consider is a comparison of human will versus spontaneous creation. The vast majority of modern life is devoted to the development of the personal will, of focusing attention toward a goal, of realising an imagined future. The irony of this is that certain individuals who attain a miraculous level of willpower have eventually discovered, often through spiritually based awakenings, that the most refined use of will is to allow it not to be the primary tool of personal guidance in one's life. To allow will to coexist with acceptance of a natural flow, a continuum of change in all things, is to realise one's most natural place in the universe, the 'grand scheme of things'.

By working within rule sets, we are implicitly required to master will-based techniques of manipulating bits of energy and matter until we achieve a desired transformation of data external to us – the 'piece'. We also will have unwittingly reinforced, and in fact refined, the will-based techniques involved, and perpetuated the tendency to use this paradigm for problem-solving and navigation through Life as a whole.

Conversely, by acknowledging that in every moment of our lives we are given a chance to 'change our minds', to 'let things go', to 'just see what happens next', our motivation for developing techniques is forever altered, and the techniques and tools themselves become different than those previously valued. While, as artists, we will continue to create encapsulated 'energy and matter things', we, ourselves, will evolve much differently than our will-based colleagues. If there is no predefined target or goal, then a new spectrum of questions must be confronted as we determine what a 'finished piece' actually is. The internal, personal affect of 'completing' a piece, using tools based not on the will, but more from an 'observer, or sensory-based' perspective, will be profoundly different. Definitions of success and failure change forever. We give ourselves infinitely numerous channels for perceiving the interconnectedness of our experience as sentient beings, rather than a set of predefined constraints with which to deal as we attempt to 'succeed' in our circumscribed creative endeavour.

3.7. Complexity and simplicity

In music theory traditions, the analysis of style or form is generally based upon a gross overemphasis on a very few arbitrary musical attributes such as metrical cycles ('the beat'), harmonic orientation ('the key or mode'), timbral palette ('orchestration') or the patterns of large-scale passages ('movements'). This approach continues to this day and is so ingrained in our thinking and listening that we tend, on the whole, to be much more like the 'unadventurous' listener than the 'adventurous' one described above.

Because we are trained to recognise what we refer to as 'style' by differentiating a very few easily perceived characteristics, we are unable to comprehend compositional structures that manipulate different conceptual materials in different ways, thereby focusing upon unfamiliar characteristics. The vast majority of people have extremely limited databases of musical experiences. This is due not to an overwhelming genetic tendency toward constraining one's musical experience. It is due to the cultural and economic structures through which musical ideas have been forced to flow for centuries. The fault lies as much with the producers as with the consumers. We have the technology to right this aesthetic injustice. The same technology may be (and is) used to enforce and deepen the musical ignorance that permeates nearly all of what most people hear each day.

The question of relative complexity is one that confronts anyone working in the arts and sciences. In our situation, this question relates to an examination of existing styles of music and understanding where each falls on a scale measuring complexity. My feeling is that all existing styles reside in a very narrow area of the scale, located much nearer to utter simplicity than to infinite complexity.

Philosophically, each artist and composer must come to terms with the question of what exactly his or her work is supposed to do for those who may encounter it in the near present or in the far future. By propagating forms that are relatively simple to experience, analyse and understand, one reinforces the notion of relatively simple intellectual, psychological and emotional behaviour. By offering relatively complex forms, one is implicitly and explicitly inferring that complex processing is required in order to adequately understand and appreciate the form that is the object of attention.

The object being focused upon by the listener will not always be musical in nature, but may be something entirely different: a personal dilemma, a political strategy, a mathematical or philosophical problem. Simplicity does indeed inform us, and gives us tools with which to deal with problem-solving. My argument is that the role of the technological arts should be to provide people with a complete spectrum of problem-solving tools, from simple to complex – a full battery of strategies for adapting to our overwhelmingly complex life. I am not convinced that we as a community of computer music artists are committed to this goal. In fact, at times I believe we are lazy, and intent on simply making it easier to create even more easily understood 'things' for the world to gawk at and identify with.

It is not a simple world, and it is sometimes a seemingly incomprehensible world. People need complex tools to deal effectively with the everyday challenge of just surviving. Why not use technological art as a mechanism for creating the metaphors needed in order to help others grasp the potential applications of complex problem-solving tools?

3.8. Constraint and boundary recognition

Computer-assisted music composition is currently pursued with one of two primary goals. Regardless of the specific technical methods and materials focused upon, the underlying motivation is to create a composition which either (i) lies within the bounds of one or more known systems of constraints (commonly called styles or traditional forms), or (ii) lies outside the accepted bounds of all constraint systems known to the composer. I argue that energy invested in the former vein of thinking is only relevant if the results

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

produced and the knowledge acquired are later reinvested as fuel for the latter motivation.

In order to create a system of thought that by its very nature precludes any repetition of, or preference for, prior stylistic constraints, we are again confronted with two approaches. Firstly, one may assume that a knowledge of all existing constraint systems is required in order to create a system that in effect does not replicate any of them. This seems to be a monumental undertaking, based on building separate conceptual models of each constraint system until all stylistic boundaries are thoroughly understood with relation to the boundaries of each other system. Alternatively, one may approach the entire problem as a question of providing a single conceptual model capable of producing compositions adhering to not only any and all constraint systems put forth in the past, but more importantly, to infinitely numerous heretofore undiscovered and unimagined ones. Though this may seem even more monumental in scope, it appears to the author to be a profoundly more appropriate use of time and energy and potentially much more significant in the 'grand scheme of things'.

3.9. Data assimilation, transformation and auralisation: musical structure as a metaphorical aid to scientific research

By using a *general model of data assimilation and transformation* (figure 2) we are able to place musical structure and music cognition in a more universal perspective. The general processes involved in the creation and perception of data structures are consistent regardless of the nature of the primary data, whether these data occur as consciously willed or unconsciously emergent event sequences. Furthermore, the data structures of one context may be mapped (interpreted) through the *general model* and be represented in the native structures of a different context. The analysis of data structures and the subsequent interpretation of this information into new musical compositions may be called *data auralisation*, though recently the term *data sonification* has taken hold. Thus, by using musical structure as a translator, the *general model* may provide a mechanism for cross-disciplinary data mapping, enhance the understanding of data within each context, and highlight relationships between contexts.

3.10. Musical structure as an experiential communications language

Music has traditionally been primarily identified as one of the arts, recognised for its evocative ability to communicate the human experience from musician to audience through the temporal vehicle of sound. As such, music has been thought of as a language as well. The analysis of the language of music and the associated teaching methods that have developed in our (human) society have become institutionalised within many social, cultural and educational traditions. Because of this, there exist many 'dialects' of music on the planet, each with its own interpretation of the nature of the language, its components, how it is to be used, what linguistic standards are required in order to communicate intelligibly, and what is not acceptable language use.

What has not been fully perceived is an understanding of how the language of music is intimately related to every other system of thought, feeling and belief that humans maintain, as well as to the natural and artificial systems that interact with us during our lives.

In expanding the contexts from which seminal data is allowed to flow into musical creation, the vocabularies and therefore the description tools available to creators of musical structures become significantly expanded as well. Hence, a greater degree of refinement and a more robust palette of syntactical constructs may be communicated between humans through the art of musical expression.

3.11. The creative process – order out of chaos or chaos out of order?

The process of creating a structure outside of one's self can be seen as an attempt to organise ideas, emotions, techniques and materials that one has collected and acquired over time, and then to combine these elements in a manner that results in a structural representation of the process that one passes through during this creative process. The structure created is essentially evidence of activity, a residual record of energy and matter having interacted, generally in an environment itself at least partially created by one or more humans.

The primal motivation to create is rooted in the tension between (i) the need to make sense of one's experience as a living, sentient being, and (ii) the need to challenge one's sense of reality by extending one's experience into previously unexplored territory. This tension is perpetually presented to us as we seek out new experience, assimilate this experience into memory, and seek again, having now modified the base of information available to us for exploration and understanding.

The *general model* is continually iterated as we continue in consciousness, searching for new data structures to evaluate and understand, to add to our stockpile of evidence that, yes, we are here . . . we are each unique . . . no one else knows exactly what we know . . . and no one can create exactly what we create!

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

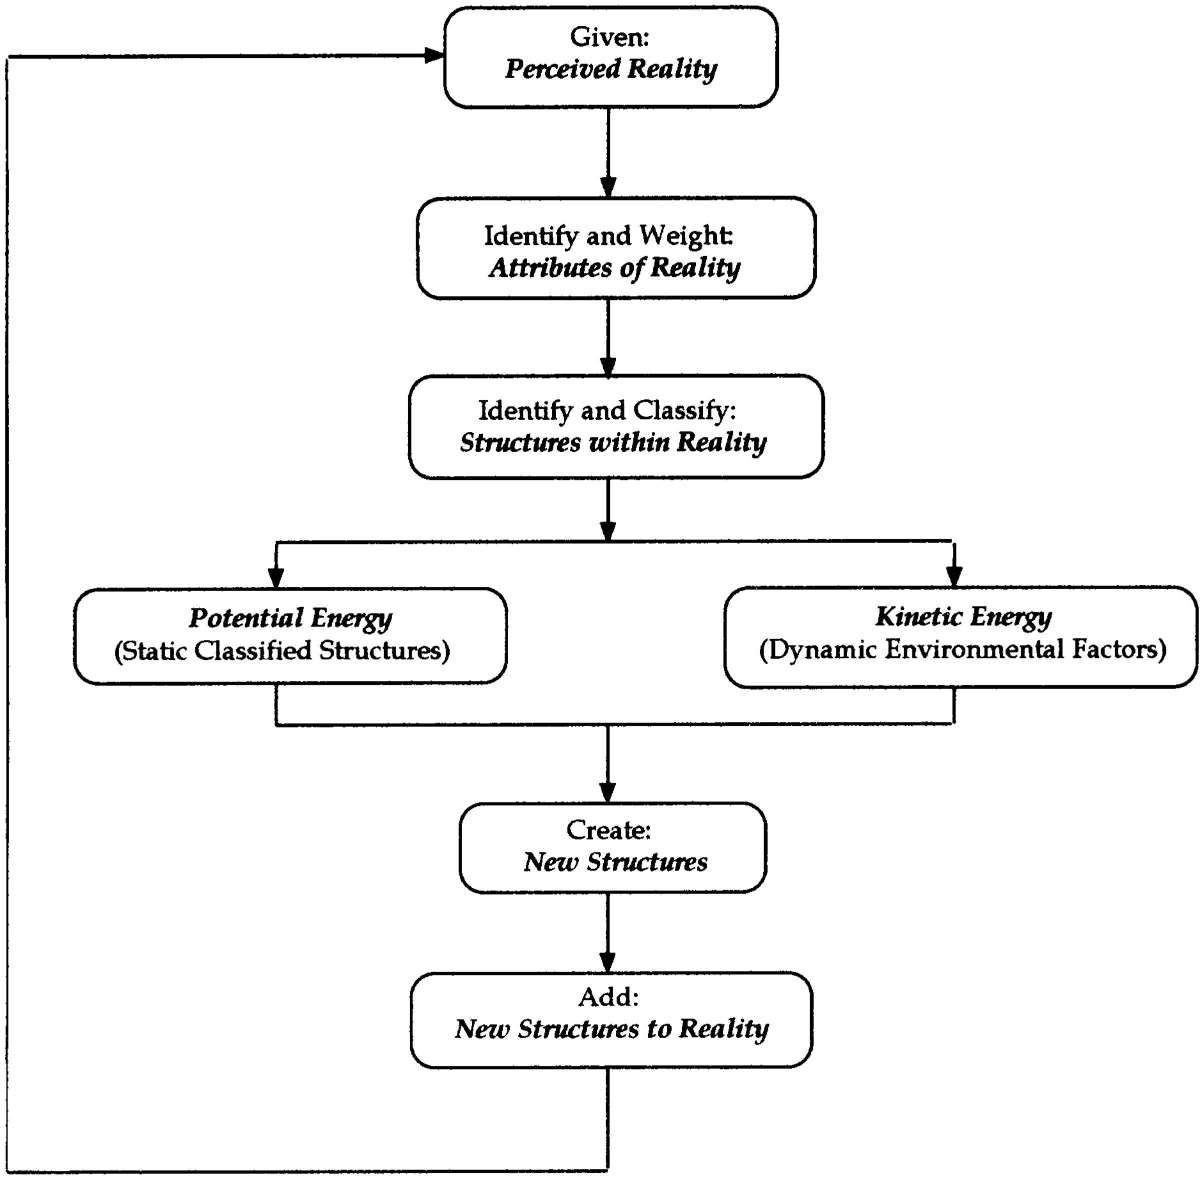

Figure 2. General model of data assimilation and transformation.

3.12. Efficiency in creative work

Redundancy in the creative arts is to be avoided above all else. It is the basic imperative of the creative impulse to explore, document, describe and interpret through one's art as much experiential territory as possible during one's lifetime. The most important skill to learn is 'tool acquisition'. The tools most highly valued are those that enable one to acquire new data, the raw materials required for expanding the boundaries of exploration and experimentation.

A tool that is designed specifically to acquire, decode and transform data from an unlimited variety of sources must therefore be considered worthy of investigation. The search for more efficient data assimilation tools is the underlying motivation for the absorption of musical styles, traditions and perspectives, leading many artists, engineers and researchers alike to the realisation that creative work is an ongoing process of data assimilation and transformation. As undiscovered data structures are created, perceived and assimilated, the base of knowledge available for the creation and perception of each subsequent structure is continually transformed.

As long as the data captured through this hierarchical iterative process is assimilated and allowed to have consequence upon subsequent creative and perceptual decisions, the amount of wasted time and energy remains very small, the level of efficiency and the potential for genuine creative excellence very high.

3.13. Contexts for complex data

Contexts for the metaphorical use of musical interpretation include the study of genetics, biological

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

populations, astrophysical exploration, biofeedback, and learning. Data from these contexts are readily mapped through the *general model* and may then be transformed into musical output data structures.

To be useful, a musical metaphor must supply a researcher with information she might otherwise not perceive, in order to better inform her of the nature of the data being examined. To do this, the metaphor requires that an inherent degree of complexity be present in the source data it analyses and manipulates. If the metaphor is valid at all, its complexity will directly mirror the complexity of the input structures. The materials available for creating a useful metaphor are defined in large part by the nature of the object it describes and the accuracy and precision of the original attributes of the object.

The above disciplines offer complex data representing physical bodies ranging in size from cellular to galactic, and behavioural processes ranging in scope from neural synapses to planetary evolution. In all cases, data auralisation may aid in the enhancement of information available for study. In addition, these fields of study all involve some form of data assimilation, so the use of a metaphorical process that also recycles information, implicitly increasing its understanding of source data structures, may likely add further refinement to the accuracy and significance of the musical metaphor.

3.14. Polarities, spectrums and domains

In the case of a purely musical context, source data structures contain data representing musical information. The procedure followed in the analysis and classification phase of processing musical source data is identical to the one followed for nonmusical contexts.

The concept of a spectrum of values within a variable domain is the basic method of quantifying seemingly intangible musical qualities so that they may be measured, evaluated and manipulated. As analysis proceeds, a design gradually unfolds that resembles a galaxy of data, each having its own spectrum or range of possible finite values. The true art to making this massive aggregation of individual values come to life is in the creation of linkages between the data, even to the level of creating different linkages based on the values of each datum at a given point in time.

When a single variable is being measured either as a discrete value or as a percentage of its potential maximum value, and quite clearly this is the case in most situations within our system, the use of *polarity* as a metaphor is perfectly valid, and becomes essential to our hierarchical system as we develop the design further. However, if a range of values is in fact representative of a summing of multiple single-variable values then we encounter a very different concept indeed. By comparing large hierarchies of linked variables we are able to place an entire hierarchy into a *spectrum* of values so as to be closest in proximity to other hierarchies that resemble it most in a composite manner. By doing this we create a spectrum whose values do not represent actual numeric quantity but relative positions within a topology of similar and dissimilar characteristics.

If we had chosen to follow the first approach described above and built a separate model for each existing musical style known to us, we would not be able to accomplish the task of accurately comparing them, since they would not be defined in identical model terms. In building a generalised model of musical structure and the compositional process, we are able to construct a circular spectrum of *domains* by placing each 'style' into a position that defines its composite distance or relationship to all other 'styles'.

4. DESIGN

A detailed component-level design specification of the MITA system is beyond the scope of this paper, and currently would include proprietary information. However, a thorough discussion of the conceptual foundation of the system is not complete without at least a few notes on design and implementation.

Clarity is the most important resource in building a model requiring such a global perspective. Scrutiny must be applied liberally in every step of an analysis and design process that is not unlike that followed when designing a computer, an operating system or a language compiler. In fact, other than the binary number system, nothing can be unequivocally accepted as a valid building block for our system. Every assumption must be questioned, examined until there is no doubt that we have overlooked nothing, that no critical piece of the puzzle has eluded us only to reappear at a later time as a deeply embedded flaw, much more difficult to eradicate than if caught earlier in the process.

As already described, familiarity plays a major role in any evaluation to be made. So, what makes a 'good' temporal structure as opposed to a 'bad' one? Simply stated, it is the level of clarity maintained during the analysis and design of the system that will in turn create the creators of these structures. For the sake of clarity, we must maintain an admittedly brutal relationship with our own personal database of musical experiences, and not allow it to influence the design disproportionately. It must be seen as only one of an infinite number of perspectives, and each conclusion reached during the design of the system must be accountable to a more visionary scope. Specifically, any design approach that is found to potentially favour one decision-making sequence over another must be discarded, and another one found that allows for a full spectrum of possibilities within whatever variable domain its focus may be.

4.1. A different type of radio dial

In our system we use the metaphor of a knob that is able to be turned 360 degrees in either direction to represent a stylistic spectrum. In fact we use two knobs! One knob represents the results of an analysis and classification of all the musical materials we have examined, placing these 'stylistic' samples into a topological domain. A knob setting shares the most characteristics with its closest neighbours, and shares the least characteristics with those located 180 degrees away. The second knob represents the results of an analysis and classification of all initial parameter settings we attach to the dynamic data variables essential for the production of new compositions. These parameters determine to a large degree how the musical material linked to the first knob is to be manipulated during a single compositional assembly process. Again, close proximity on this knob equates to a similarity in characteristics over multiple variables, and 180 degrees of polarity equates to very little sharing of characteristics. These two knobs are able to be manipulated independently, yielding a tremendous number of parameter/music material combinations. We refer to these knob settings as index/parameter pairs. The index term refers to the index matrix derived during the classification process, and the parameter term refers to the initial parameter settings.

Comparing the index/parameter knobs to AM/FM radio dials we are perhaps able to illustrate the profound difference between navigation within our system of musical creation and the world of contemporary musical consumption. Each radio station has its frequency domain, as each position on our two knobs have their classification or style domain. As we turn the radio dial we encounter different styles, musical and editorial, but all of the music we find is prerecorded. Sometimes a live performance is encountered, but these too are 'preprogrammed'. Alternatively, as we turn our two knobs to new positions we find that we are using a set of data that then interacts with us and itself to create music that has never been heard before, and will almost definitely never be recreated. The positions of the knobs do have a stylistic relevance, but the music pursuant from these positions is open to the influence of the moment.

4.2. Macroaesthetics, synthetic genetics, and the natural world as control mechanisms

In order to ensure that our general model does in fact fully represent a robust compositional process capable of limitless variations, we must instill mechanisms that impact the process dynamically as a decision-making hierarchy unfolds. A critical aspect of these mechanisms is that they interact with one another and with some human activity. It is only then that they have a physical link with us and become an extension of ourselves. A matrix of self-organising knowledge domains that continually interact appears to be an appropriate modelling platform for achieving an underlying level of the intangible cohesion, elegance and complexity that so often signify work described as 'beyond human capability', 'pure genius', 'profound', or simply 'very interesting'.

I have chosen several self-perpetuating models, including genetic algorithms, chaotic systems and Markov chains. The MITA population of abstract organisms is patterned after biological genetics. It is significant that though genetic algorithms are traditionally employed to solve problems, the role of our population is not to provide answers to problems. Instead, we monitor the population's activity and use its process rather than its eventual conclusions to affect our assembly process.

Activity in the external environment is captured through use of a sensing device, e.g. a video camera. Periodic 'snapshots' or 'samplings' of the environment are taken. These are analysed and compared in order to discern the relative level and type of activity occurring outside the digital world of our computer. The data derived from our internal and external measurement devices interact at critical junctures within the assembly processing path, and have a great impact on how the original musical database is manipulated.

One example of control mechanism interaction involves measuring the adaptability of our genetic population to activity in the external world. Presence of change in the external environment triggers decreased adaptability in the genetic domain. Subsequent musical event choices reflect the increased instability in the external environment, creating a less stable, inconsistent and more unpredictable musical form. Absence of change in the external environment triggers increased adaptability in the genetic domain. As the population becomes better at predicting the behaviour of a static world, the eventual result is more predictable choice patterns, and therefore a more stable, consistent and predictable compositional structure.

As described earlier, an important design component involves implementing the concept of hierarchical linkages within our musical database, or more precisely, within the variables that define our perspective upon this musical database. Several high-level control variables representing conceptual characteristics such as abstraction, density and repetition, modify themselves during the assembly process. Each of these variables is linked to a hierarchy of variables whose values directly impact the characteristic represented. As the control variables change so do their hierarchies and, of course, the fluctuation of each of these control variables is linked to some activity within the other self-perpetuating models.

4.3. Deconstruction and classification

The preprocessing of musical performance data, whether improvised, composed or computer generated, consists of analysing and encoding the mathematical profile existing within a musical expression, and formatting the resulting analysis so as to be usable in compositional tasks.

Linguistically, the numerical source data of a musical passage can be seen as an evolving language and vocabulary, creating an inherent functional classification unique to this musical context. The classification and the attendant model of auditory perception can be isolated using any number of classification methods currently being explored, such as neural nets, rough sets, fuzzy logic, SG nets, or probabilistic systems of Bayesian analysis. When statistical analysis is applied to a musical passage, the intentions and hypotheses of the researcher and the significance of the results can best be determined by an examination of the modelling techniques employed. There is currently much emphasis placed on the deconstruction of easily identified traditional 'musical styles' with the implicit intent being to replicate the style in subsequent automated compositions. In modelling data for an analysis aimed toward these ends, it is essential that the attributes of the model conform to a pre-existing 'definition' of the style being analysed. In other words, the analysis has already taken place in the mind of the researcher prior to any software modelling, and the analysis is of a 'style', not of Music. This situation guarantees a very limited flexibility in not only the resulting compositional structures but to the applicability of input data as well. Ironically, these limitations are viewed by many to be evidence of successful music composition software, rather than reformulations of well-documented arbitrary rule systems reinforced through cultural traditions.

In the MITA system, great care is taken to circumvent any perceptual bias or limitation imposed by innate, cultural or habitual memory and expectation patterns with regard to musical structure and expression. Rather than the identification and segregation of musical traditions and styles, the intent is to create a generalised model of musical form, to break down the lines of categorisation and deal with music as a single organic substance. This is accomplished by providing a capability that allows for customised dynamic models to be constructed from a set of dynamically defined individual model attributes. The attributes themselves are developed by taking the mathematical information present in the source data, creating hierarchical relationships within the attribute space, and yielding an attribute set of potentially infinite size.

Classification is able to be applied to one or more musical event types, such as pitch, duration, volume, etc. We deconstruct the source data into fixed-length arrays that portray the values of the specified event types over a finite length of time or for a number of events. This process of deconstruction can continue for as many iterations as is desired for numerous combinations of model attributes. Each set of fixed-length arrays is treated as a single database and used for a distinct classification cycle. The performance source data may thus be analysed from several perspectives based on differing quantities and combinations of attributes. The resulting analysis is translated into numerous data structures that reference the original performance datastream while retaining the accumulated classification data. In effect, this iterative deconstruction and classification approach allows for precisely classified phrases of varying length and, more profoundly, each numerical component of each event within the entire event stream to be made available as data-primitives for use in assembly decision-making.

4.4. Assembly

The task of combining and assembling discrete musical/mathematical materials into integrated and coherent musical pieces is what people generally refer to as composition. MITA assembly processing is extremely complex, as you would assume, but in general is based on several nested loops that together produce a classic straight-ahead decision-making path. Each loop is begun with a refresh of certain data structures that are critical to decision-making within that loop, creating a consistent context that informs the logic as it proceeds in its choices of branching direction and of data values. This continually refreshed context is derived from (i) the intelligence embedded in the classified performance data that also acts as primary musical source material, (ii) the sequence of decisions already made during the assembly, and (iii) the current state of the compositional structure itself.

Additional leverage at any given point during the construction process is provided by the continual monitoring of several simultaneously interacting periodic and aperiodic control mechanisms that act as loop interrupts. They include (i) an abstract hierarchical environment of mathematical organisms that evolves over time, (ii) a set of macro-parameters, whose values change over time, describing the high-level structural underpinnings that implicitly affect all

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

aesthetic decision-making, and (iii) periodic samplings of an external event system that reflect movement or change of some sort within a physical or abstract universe.

The current system is MIDI based, and provides for continuous controller and system exclusive event processing as well as for basic note and velocity events. The fully assembled composition is performed through sound source equipment and captured in a Standard MIDI File (SMF). For subsequent reference purposes, an audit log is produced to reflect in detail all the decision-making activity that has taken place during the assembly process.

MITA v0.71 software consists of approximately 25,000 lines of source code written by the author using the Forth-based languages HMSL and H4th (Frog Peak Music 1990–7) on the Macintosh platform.

4.5. Compositional results

Compositions produced by the system include *Daydreams of an Orange Cat*, created on October 20th 1993, using software version 0.60. This piece was premiered at Clark University, Worcester MA as part of the Studio Music Concert Series on March 10th 1994.

A more recent piece, *Run The Mountain*, was created on January 10th 1997 using software version 0.651, and was premiered at Aristotle University, Thessaloniki, Greece as part of the 1997 International Computer Music Conference on September 27th 1997.

5. CONCLUSIONS

The significance of creating and perceiving infinitely numerous unique forms should be apparent to anyone seriously involved with arts and technology. The enhancement of a person's life through the broadening and deepening of perception is the common ground binding the pursuits of intellectual education, physiological evolution and spiritual enlightenment. The theoretical model presented here draws upon these foundations of human existence in an attempt to integrate and synthesise the human experience into a music-based metaphorical tool. In a very true sense, this entire general model of composition is a self-perpetuating artificial organism of immense complexity. Though initially designed to manipulate musical data, the system hopefully contains genuine significance for non-musical applications as well.

The Music In The Air system architecture provides for an infinitely variable and complex numerical/musical database for performance and composition, based in part on improvised and composed material and in part upon event-streams outside the realm of musical thought, giving the user streamlined control mechanisms for the simultaneous manipulation of all process levels.

At each level of the potentially dense decision-making process there are numerous options as to what to do next, and a unique combination of available data that contribute to the decision to be made. All during the construction process, regardless of how much activity takes place within the control mechanisms and their associated internal data mappings, the actual musical materials being manipulated remain the result of human/machine musical/non-musical activity as seen from the perspective of a rigorous statistical analysis.

Research in the fields of computer science and mathematics yield quantifiable results that are able to be judged with rigour on a basis of accuracy or optimisation. Experimental research in the field of computer music yields results in the form of musical compositions and performances. The theoretical model and complex software system described here are not in themselves results that are able to be quantifiably judged as to their relative merit. Regardless of the technical rigour with which a music system is designed and implemented, the system itself is invariably judged, or at least should be, on the basis of the musical substance it produces. This demands a tremendous aesthetic flexibility on the part of those passing judgement. We must measure the conceptual perspective, the technical implementation, and the musical results, based in part on their own terms, and in part against historical and contemporary research spanning numerous technical and artistic disciplines.

. . . Infinitely numerous forms exist,
and potentially exist,
as unique amalgams
of seemingly random sequences of processes
acting upon fragmented aspects of sonic matter. . .

In pure silence we hear nothing, but we listen just the same, to our thoughts and senses existing in perpetual activity. As silence retreats, the sonic entity in our external world enters in and becomes part of the constant exchange inside us. . . . Within the realm of music lies the potential for transmuting the holistic and transcendental nature of human experience into countless forms and processes that interact to produce infinitely varied sonic structures.

. . . If the music speaks with equal clarity
to both a wide-eyed child of two and
to an adult seeker of intellectual and sensual challenge,
then the music is,
in fact,
something of value. . .

REFERENCES

Allik, K., and Mulder, R. C. F. 1993. SKYHARP: an interactive electroacoustic soundscape installation. *Proc. of the Fourth Biennial Connecticut College Arts & Technology Symp.*, pp. 190–8. Center for Arts & Technology at Connecticut College.

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press

Alphonce, B. H. 1980. Music analysis by computer – a field for theory formation. *Computer Music Journal* **4**(2): 26–35.

Ames, C. 1983. Stylistic automata in *Gradient*. *Computer Music Journal* **7**(4): 45–56.

Ames, C. 1991. Formalism and musical composition. *Proc. of the Third Biennial Connecticut College Arts & Technology Symp.,* pp. 167–79. Center for Arts & Technology at Connecticut College.

Ames, C. 1992. Quantifying musical merit. *Interface* **21**: 53–93.

Bauer, W., and Foss, B. 1992. GAMS: an integrated media controller system. *Computer Music Journal* **16**(1): 19–24.

Beyls, P. E. 1989. The musical universe of cellular automata. *Proc. 1989 ICMC*, pp. 34–41. International Computer Music Association.

Beyls, P. E. 1991a. Self-organising control structures using multiple cellular automata. *Proc. 1991 ICMC*, pp. 254–7. International Computer Music Association.

Beyls, P. E. 1991b. Chaos and creativity. *Proc. of the Third Biennial Connecticut College Arts & Technology Symp.,* pp. 157–65. Center for Arts & Technology at Connecticut College.

Bharucha, J. J., and Todd, P. M. 1989. Modeling the perception of tonal structure with neural nets. *Computer Music Journal* **13**(4): 44–53.

Bidlack, R. 1992. Chaotic systems as simple (but complex) compositional algorithms. *Computer Music Journal* **16**(3): 33–47.

Brubaker, D. 1993. An accelerated kernel for fuzzy systems. *AI Expert* **8**(3): 38–43.

Burk, P., Polansky, L., and Rosenboom, D. 1990. *HMSL: Hierarchical Music Specification Language – Version 4.0*. Frog Peak Music.

Buxton, W. 1977. A composer's introduction to computer music. *Interface* **6**: 57–72.

Buxton, W. 1978. Design issues in the foundation of a computer-based tool for music composition. *Technical Report CSRG-97*. Computer Systems Research Group, University of Toronto.

Caudill, M. 1993. The possibilities of probabilities. *AI Expert* **8**(3): 28–31.

Chabot, X. 1990. Gesture interfaces and a software toolkit for performance with electronics. *Computer Music Journal* **14**(2): 15–27.

Chadabe, J. 1984. Interactive composing: an overview. *Computer Music Journal* **8**(1): 22–7.

Cohn, D., Atlas, L., and Ladner, R. 1993. Improving generalisation with self-directed learning. *Machine Learning*.

Cope, D. 1992. Computer modeling of musical intelligence in EMI. *Computer Music Journal* **16**(2): 69–83.

Dawkins, R. 1987. *The Blind Watchmaker*. New York: W. W. Norton.

Dewdney, A. K. 1988. *The Armchair Universe*. New York: W. H. Freeman.

Dodge, C., and Jerse, T. 1985. *Computer Music: Synthesis, Composition and Performance*. New York: Schirmer Books.

Dyer, C. 1993. Three-dimensional machine vision. *Computer Science Department Symp. at Dartmouth College*.

Fry, C. 1980. Computer improvisation. *Computer Music Journal* **4**(3): 48–58.

Gogins, M. 1991. Iterated functions systems music. *Computer Music Journal* **15**(1): 40–8.

Goldspiel, J. M., and Squyres, S. W. 1991. Ancient aqueous sedimentation on Mars. *Icarus 89*, pp. 392–410. Orlando, FL: Academic Press, Inc.

Hanninen, D. 1996. A general theory for context-sensitive music analysis. *Society for Music Theory 1996 National Conf.* Baton Rouge, LA.

Horner, A., and Goldberg, D. E. 1991. Genetic algorithms and computer-assisted music composition. *Proc. 1991 ICMC*, pp. 479–82. International Computer Music Association.

Koenig, M. G. 1983. Aesthetic integration of computer-composed scores. *Computer Music Journal* **7**(4): 27–32.

Langston, P. 1989. Six techniques for algorithmic music composition (extended abstract). *Proc. 1989 ICMC*, pp. 164–7. International Computer Music Association.

Laske, O. E. 1972. On musical strategies with a view to a generative theory of music. *Interface* **1**: 111.

Laske, O. E. 1975. Toward a theory of musical cognition. *Interface* **4**: 147.

Laske, O. E. 1980. Toward an explicit cognitive theory of musical listening. *Computer Music Journal* **4**(2): 71–83.

Laske, O. E. 1988. Introduction to cognitive musicology. *Computer Music Journal* **12**(1): 43–57.

Laske, O. E., Balaban, M., and Ebcioglu, K. (eds.) 1991. *Understanding Music with AI*. Menlo Park, CA: AAAI Press.

Low, A. 1991. *Introductory Computer Vision and Image Processing*. London: McGraw-Hill.

Loy, G. 1991. Connectionism and musiconomy. *Proc. 1991 ICMC*, pp. 364–74. International Computer Music Association.

Newcomb, R. S. 1980. *ARSYN I: the arbitrary and synthetic nature of sequential logic*. Unpublished monograph.

Newcomb, R. S. 1981. *ARSYN II: the arbitrary and synthetic nature of sequential logic*. Unpublished monograph.

Oppenheim, D. 1989. DMIX: an environment for composition. *Proc. 1989 ICMC*, pp. 226–33. International Computer Music Association.

Owens, T. 1974. *Charlie Parker: Techniques of Improvisation, Vol. I & II.* UCLA PhD Dissertation. University Microfilms International.

Polansky, L. 1987. Morphological metrics: an introduction to a theory of formal distances. *Proc. 1987 ICMC*, pp. 197–205. International Computer Music Association.

Polansky, L., and McKinney, M. 1991. Morphological mutation functions. *Proc. 1991 ICMC*, pp. 234–41. International Computer Music Association.

Pope, S. T. 1992. The Interim DynaPiano: an integrated computer tool and instrument for composers. *Computer Music Journal* **16**(3): 73–91.

Pressing, J. 1988. Non-linear maps as generators of musical design. *Computer Music Journal* **12**(2): 35–46.

Roads, C., and Stawn, J. (eds.) 1987. *Foundations of Computer Music*. Cambridge, MA: MIT Press.

Rosenboom, D. 1990. The performing brain. *Computer Music Journal* **14**(1): 48–66.

Rowe, R. 1992. Machine listening and composing with Cypher. *Computer Music Journal* **16**(1): 43–63.

Sagan, C. 1980. *Cosmos*. New York: Random House.
Sagan, C. 1985. *Contact*. New York: Simon & Schuster.
Stone, P., and Morgan, N. 1993. *Symbolic Composer*. Commercial marketing brochure, Tonality Systems.
Szladow, A., and Ziarko, W. 1993. Rough sets: working with imperfect data. *AI Expert* **8**(7): 36–41.
Taylor, W., Stutz, J., *et al*. 1992. Bayesian classification scheme. *NASA Technical Support Package ARC-12797.* NASA.
Todd, P. M. 1989. A connectionist approach to algorithmic composition. *Computer Music Journal* **13**(4): 27–43. Cambridge, MA: MIT Press.
Tompkins, P., and Bird, C. 1973. *The Secret Life of Plants*. New York: Avon Books.
Truax, B. 1978. *Handbook for Acoustic Ecology*. Vancouver, BC: A.R.C.
Truax, B. 1980. The inverse relation between generality and strength in computer music programs. *Interface* **9**: 49.
Truax, B. 1992. Musical creativity and complexity at the threshold of the 21st century. *Interface* **21**: 29–41.
Vivid Effects, Inc. 1992. *Mandala Interactive Software*. Commercial marketing brochure.
Widmer, G. 1992. Qualitative perception modeling and intelligent musical learning. *Computer Musical Journal* **16**(2): 51–68.
Winsor, P. 1991. PAT-PROC: an interactive, pattern-process, algorithmic composition program. *Proc. 1991 ICMC*, pp. 114–17. International Computer Music Association.
Xenakis, I. 1971. *Formalized Music*. Bloomington: Indiana University Press.

https://doi.org/10.1017/S1355771898009121 Published online by Cambridge University Press