Preface

Since the early 1950s, artificial intelligence has led a very public life, gaining visibility in the press, novels, television, motion pictures, and predictionists' visions of the future. Virtually hundreds of scholarly and not so scholarly books and articles have appeared, covering the history, evolution, successes, failures, and potentials of artificial intelligence. Marketing claims ranging from intelligent word-processing software to intelligent toasters continue to emerge almost daily. It is difficult to separate the actual from the hype and even more difficult to define intelligence, so nonchalantly is the term wielded. Amid all of this attention to artificial intelligence, the term "artificial creativity" has barely surfaced. It may be that we assume that creativity requires intelligence and that the latter subsumes the former. Whatever the reasons, however, artificial creativity has followed a much more sedate path during these past fifty years, and the few of us who labor to understand and model it have had to work in relative anonymity—possibly to our advantage, since the resultant low expectations have often matched our achievements.

As I hope to prove, "artificial creativity" is too broad a term to cover the study of the various ways that humans invent new art. Language, music, visual art, dance, and so on all develop in unique ways by utilizing different combinations of our creative abilities. In this book I will concentrate primarily on artificial musical creativity—the organization of sounds and silences that do not possess agreed-upon meanings. As a close relative of both poetry and dance, music offers certain ineffable qualities not found in any other of the arts. Programming or other ways of simulating creative processes in music require an understanding of music's unique modes of expression and contexts of performance and perception. In short, even the field of musical creativity is extraordinarily broad and complex. My hope, therefore, is modest: to model, using computers, the more salient features of musical creativity in order to gain a better understanding of a few of the processes that define and distinguish it from creativity in the other arts.

I will also attempt to grapple with questions concerning the definitions of creativity, whether or not computer programs can effectively model creativity, and whether or not computer programs can in fact create. This grappling will include the views of many computer scientists, musicians, scholars, and psychologists who have seriously engaged this subject, including such distinguished individuals as Douglas Hofstadter, Margaret Boden, Selmer Bringsjord, David Ferrucci, Karl Pfenninger, Valerie Shubik, John Dacey, Kathleen Lennon, and many others. Having been a professional composer for over fifty years and a programmer for almost thirty years, I hope to add my voice to theirs in some small way, especially as creativity applies to music. Interestingly, my views prove quite different from those I have encountered. For example, I will contend that computer programs *can* create. I will further contend that

those who do not believe this have probably defined creativity so narrowly that humans could not be said to create.

Whenever possible in this book, I have included the thoughts of professional musicians, psychologists, philosophers, computer scientists, and cognitionists whose work complements or poses the greatest challenges to my own ideas. My apologies to those whose work may seem relevant, but to which I have not referred because of space limitations.

I do not intend this book as a comprehensive guide for related research on the topics covered. Such a book would require several thousand pages, with the thread of how the models form a cohesive whole possibly lost in the diversity of ideas. In place of comprehensive coverage, I have structured this book in three major sections—"Background and Principles"; "Experimental Models of Musical Creativity"; and "An Integrated Model of Musical Creativity"—whose titles I believe roughly describe their contents. In these sections, I make a case for musical creativity resulting from a process called *inductive association*. In brief, inductive association is a more narrowly defined version of *free association*, or the shedding of deductive reasoning for more intuitive processes.

Each chapter of this book begins with a simple principle, one that I attempt to prove as that chapter continues. These principles are followed by illustrative vignettes (personal and otherwise) apropos to that chapter's subject. In this way, I hope to create a friendly narrative while maintaining a scholarly focus. Many chapters also contain descriptions of computer programs designed to demonstrate that chapter's focus on the complexities of musical creativity (see appendix D for a complete listing of the programs described in this book and available on my Web site).

Part I of this book—"Background and Principles"—provides the historical and definitional foundations of creativity, so important for understanding what we mean when we use that word. To this end, chapter 1 presents a number of different perspectives on music and meaning. It continues by defining creativity—for this book, at least—and the various ancillary terms that often accompany that word. The chapter concludes with a discussion of the role of originality in creativity. Chapter 2 presents some thoughts on the origins of creativity studies, followed by a brief history of creativity research in general and of musical creativity research in particular. It concludes with descriptions of various proposals to test for creativity. Chapter 3 begins with a detailed analysis of randomness and how it differs from creativity. It continues with a summary of computer program types that in some way may model music creativity.

Part II of this book—"Experimental Models of Musical Creativity"—describes a number of possible models for computationally imitating human creativity. Chapter

4 outlines the basic principles of recombinance and pattern matching, two foundational principles in my work with computers and music. The chapter concludes with a discussion of the problems that performance creates for computation—particularly recombinant computation. Chapter 5 describes how allusions contribute to musical creativity, and concludes with the description of a program that analyzes music for its references to other music and possible ways in which these references might be interpreted. Chapter 6 explains the role that learning plays in the creative process, then continues with a discussion of inference and how it can enhance creativity. The chapter concludes with examples of how analogy contributes to the creative process. Chapter 7 presents some of the ways in which composers build musical expectations and then either fulfill them or surprise listeners. It then discusses musical hierarchy and how computer programs can incorporate the analytical tools necessary to meld hierarchy into their creative processes. Chapter 8 describes the role of influence on creativity. It also includes a description of a program—called a *spider*—capable of independently connecting to the Internet and downloading certain types of files.

Part III of this book—"An Integrated Model of Musical Creativity"—presents an inductive-association computational process that can solve problems and produce music creatively. Chapter 9 defines association networks and explains how such networks can, without being programmed to do so, respond effectively to input. It next explores how inductive association can produce interesting, creative, and often insightful output, as opposed to the complex but straightforward noninductive output of association networks. Chapter 10 applies the principles of association networks to music, both in the form of brief exchanges and in terms of longer, more formal musical compositions. Chapter 11 discusses a number of possible combinations of the processes discussed in this book, ultimately favoring an integrative model. Chapter 12 then presents a number of the aesthetic difficulties encountered when computationally modeling creativity. This chapter continues by exploring many of the contradictions that arise when building programs capable of musical creativity. I have also included here my rationale for discontinuing my work with historical musical styles using Experiments in Musical Intelligence. The book concludes with a description of various possible futures and the role in them that machine creativity might play.

I struggled long and hard over how to present the material in this book. In fact, the first draft did not look anything like this final version. This first draft described the integrated model of creativity currently found in chapters 9 through 11 first, and then detailed the ways in which this model solved the various problems associated with studying creativity using computers. However, I finally opted for an approach that initially presents the foundations of my research and then describes various

models that solve selected aspects of creativity but, because they do not solve them all, fail to succeed in truly modeling creativity. I then follow these failed models with a description of an integrated model, the one with which I began my first draft. This latter approach—leaving the solution until near the end of the book—may frustrate some readers who do not wish to wait so long for a resolution of the various problems encountered in such studies. While I apologize to those individuals for their wait, I truly feel that the time spent reviewing the fundamentals, the problems, and the various attempted, but failed, solutions makes for a truer understanding and appreciation of the model that I feel ultimately succeeds.

While the majority of examples I provide in this book relate to music, I have also included a number of references to language, puzzles (games), and astronomy to clarify what might otherwise confuse readers were I initially to give examples in music. I have also chosen these nonmusical examples to help broaden the points I make to suggest that while my subject is musical creativity, the principles can often be extended to creativity in other fields.

I have limited the scope of study for this book to classical music. Other types of music might work equally well. However, classical music provides a comprehensive range of styles over a significant historical period. As well, my own background consists almost exclusively of classical music, and hence I lack the expertise to engage other styles. Readers may offset this shortcoming by applying the techniques defined and described here—as well as the programs offered—to whatever type of music they know best.

While readers should have some experience with music notation in order to understand many of the examples in this book, no knowledge of computer programming is required. Likewise, readers need not own specialized computer hardware or software to follow the text. I have written this book using nontechnical terminology and in a style that I hope will appeal to the general layperson with an interest in music and creativity. The narrative as well as the music should be understandable to anyone who is curious about the role computer programs can play in modeling creativity.

This book includes examples of output of the various computer programs discussed, and these programs, along with MP3 versions of all of the musical examples, are available at my Web site, arts.ucsc.edu/faculty/cope/software/cmmc. The programs are written in Common Lisp with many having two versions available: (1) Macintosh platform and (2) any platform that supports Common Lisp. To ensure that these latter programs will perform in different environments requires that I omit all MIDI (Musical Instrument Digital Interface) and graphical user interface (GUI) code, which are platform dependent. I have included full documentation and operating instructions for each program along with the code.

As time permits, I will also include code for graphics and music playback for various other platforms on my Web site. Note, however, that if history proves to be a good prognosticator of the future—in this case meaning that as soon as I write new platform-dependent programs, system hardware or software changes make my code obsolete—I can guarantee only that the cross-platform code will function without problems. In short, the software for the programs in this book, while very helpful in demonstrating the principles of each chapter and in clarifying the model of creativity that I propose, is not critical to the understanding of the material presented.

The computer output presented in this book was produced using programs similar to the programs found on my Web site, but which occasionally contain more elaborate code. Many individuals purchase books such as this for the programs that accompany them. In fact, these individuals often treat such books as *user manuals* for their associated software. *Computer Models of Musical Creativity* does not fulfill the role of a user manual. In fact, if anything, the opposite is true: the relevant software that appears on my Web site should be considered a bonus—not a requisite—for readers to better understand the principles and ideas that the book contains.

I wish to thank the many individuals who have contributed to my ongoing studies of musical creativity, particularly Douglas Hofstadter, who continues to lecture and otherwise publicly discuss my work. Thanks go as well to Patricio da Silva and Jennifer Logan for publishing Experiments in Musical Intelligence's music, which was unavailable for some twenty years until their company, Spectrum Press, made it available. I also wish to thank the Center for Computer Assisted Research in the Humanities (CCARH), especially Eleanor Selfridge-Field and Walter Hewlett. Without the moral support and advice from colleagues such as these, this book could not have been completed.

Computer Models of Musical Creativity describes and demonstrates some of the ways in which I believe we create music. I have been cautious in my estimation of the value of new ideas, preferring to believe instead that newness is often merely an enlightenment of older, but possibly lesser-known, ideas.

In all of my work, there is a strong reliance on musical analysis rather than on the development of formulas. Often this analysis centers, in part, on venerable masterworks. The associations, allusions, and so on, I discuss in this book existed in these works long before my research began. While I appreciate whatever importance readers may ascribe to my ideas, I direct their attention specifically to the results of this research, for it is from the understanding of these results that any true enlightenment my work has to offer will come.

David Cope